

pMatlab v2.0 Function Reference

Hahn Kim, Nadya Travinin, Jeremy Kepner
{hgk, nt, kepner}@ll.mit.edu

Table of Contents

pMatlab	6
pRUN	6
Np.....	6
Pid	7
pMATLAB *Deprecated*	7
pMatlab_Init *Deprecated*	7
pMatlab_Finalize *Deprecated*	8
pMatlab_ver	8
MPI_Abort	8
MatMPI_Delete_all.....	9
MPI_Run *Deprecated*	9
SendMsg	10
RecvMsg	10
Distributed matrices and matrix manipulation	11
Elementary distributed matrices	11
map/map.....	11
map/zeros	15
map/ones	16
map/rand	17
Basic array information.....	18
dmat/size	18
dmat/ndims.....	18
dmat/display	19
map/display	19
Distributed array information	20
dmat/global_block_range.....	20
dmat/global_block_ranges	21

dmat/global_ind	22
dmat/global_inds	23
dmat/global_range	24
dmat/global_ranges	26
map/inmap	28
Matrix manipulation	29
dmat/find	29
Distributed matrix manipulation	30
dmat/agg	30
dmat/agg_all	31
dmat/local	31
dmat/put_local	32
dmat/synch	33
dmat/sync2	33
.....	34
remap	34
dmat/subsasgn	34
dmat/subsref	36
map/subsasgn	37
map/subsref	37
transpose_grid	38
Elementary math functions	39
Trigonometric	39
dmat/sin, dmat/cos, dmat/tan, dmat/sec, dmat/csc, dmat/cot	39
dmat/sind, dmat/cosd, dmat/tand, dmat/secd, dmat/cscd, dmat/cotd	40
dmat/sinh, dmat/cosh, dmat/tanh, dmat/sech, dmat/csch, dmat/coth	41
dmat/asin, dmat/acos, dmat/atan, dmat/asec, dmat/acsc, dmat/acot	42
dmat/asind, dmat/acosd, dmat/atand, dmat/asecd, dmat/acscd, dmat/acotd	43
dmat/asinh, dmat/acosh, dmat/atanh, dmat/asech, dmat/acsch, dmat/acoth	44
Exponential	45
dmat/exp	45
dmat/expm1	45

dmat/log	45
dmat/log1p	46
dmat/log10	46
dmat/log2	46
dmat/pow2.....	47
dmat/realpow.....	47
dmat/reallog	47
dmat/realsqrt	48
dmat/sqrt	48
Complex	49
dmat/abs	49
dmat/angle.....	49
dmat/complex.....	49
dmat/conj.....	50
dmat/imag	50
dmat/real	50
Rounding and remainder.....	51
dmat/fix	51
dmat/floor.....	51
dmat/ceil.....	51
dmat/round	52
dmat/sign.....	52
Operators and special characters	53
dmat/plus.....	53
dmat/minus.....	53
dmat/mtimes.....	54
dmat/times.....	54
dmat/power	55
dmat/ldivide	55
dmat/rdivide	56
dmat/eq.....	56
map/eq.....	57

dmat/gt	57
map/ne	57
dmat/transpose	58
dmat/transpose	58
Sparse matrices	59
map/sparse.....	59
map/spalloc	60
dmat/sparse	60
Data analysis and Fourier transforms	61
dmat/conv2.....	61
dmat/fft.....	61
Data types and structures.....	63
dmat/double.....	63
dmat/single.....	63
dmat/uint8, dmat/uint16, dmat/uint32, dmat/uint64	64
dmat/int8, dmat/int16, dmat/int32, dmat/int64	64
Signal Processing Toolbox.....	65
dmat/dct.....	65
dmat/idct	65
Index.....	66

Introduction

This document is meant to be a reference for functions that are of use to pMatlab *users*, i.e. application developers. There are a number of additional functions included in pMatlab, but those functions are used internally by pMatlab and should **not** be called by pMatlab applications.

The functions described in this reference are divided into sections that approximately match Mathwork's own categorization of MATLAB functions. Most sections describe overloaded MATLAB functions; some sections contain additional functions that are unique to pMatlab, but are related to the overloaded MATLAB functions in that section.

- **pMatlab** describes general pMatlab functions required by all pMatlab applications.
- **Distributed matrices and matrix manipulation** describes functions related to creating and obtaining information about distributed matrices. Because this section contains a large number of overloaded MATLAB functions and a number of new pMatlab functions, it is further divided into subsections.
 - **Elementary distributed matrices** describes functions related to the creation of distributed matrices.
 - **Basic array information** and **Distributed array information** describe functions that obtain information about a distributed matrix.
 - **Matrix manipulation** and **Distributed matrix manipulation** describe functions used to manipulate distributed matrices.
- **Elementary math functions, Operators and special characters, Sparse matrices, and Data analysis and Fourier transforms** describe functions and operators that have been overloaded in pMatlab.

Most functions are *class* functions. Consequently, the names of these functions have been prefaced with the name of class they are a part of. For example, the `fft` function for the `dmat` class is listed as `dmat/fft`. Help for every function can be obtained from the MATLAB command prompt by running `help class/function`. For example, to get the help documentation for the `fft` function overloaded for `dmat`, run:

```
help dmat/fft.
```

For some overloaded functions, it may be useful to refer to the help documentation for the original MATLAB function by running `help function` at the MATLAB command prompt. To get the help documentation for the original MATLAB `fft` function, run:

```
help fft.
```

pMatlab

pRUN

Used to run a pMatlab v2 program.

Syntax

```
eval(pRUN(mfile, Ncpus, cpus))
```

Description

`mfile` is a string that contains the name of the pMatlab program to be launched, without the `.m` suffix.

`Ncpus` is an integer that specifies the number of processors to launch `mfile` onto.

`cpus` specifies what machines to launch `mfile` onto:

- `cpus = {};` Run all MATLAB processes on the local machine.
 - `cpus = {'machine1' 'machine2' ...};` Specify names of machines on which to run. To run interactively, `machine1` must be the name of local machine.
 - `cpus = {'machine1:dir1' 'machine2:dir2' ...};` Specify machines names and which directory to use for communication on each machine. Directories must be visible to both machines, i.e. crossmounted. Directories should be located on the local disk of their respective machines.
 - `cpus = {'machine1:type' 'machine2:type'};` Specify machine names and the type of each machine. `type` can be either `'unix'` or `'pc'`. Default is `'unix'` (can be changed in `MatMPI_Comm_settings.m`)
 - `cpus = {'machine1:type:dir1' 'machine2:type:dir2'};` Specify machine names, communication directories, and the type of each machine.
-

Np

Returns the total number of Matlab instances currently running (i.e. `Ncpus`).

Syntax

```
Np
```

Description

`Np` is an accessor function to the `pMATLAB` global variable and returns the total number of Matlab instances launched by `pRUN`.

Pid

Returns the `Pid` of the current instance of Matlab.

Syntax

`Pid`

Description

`Pid` is an accessor function to the `pMATLAB` global variable and returns the current instance of Matlab. `Pid` ranges from 0 to `Np`-1.

pMATLAB *Deprecated*

Data structure created by `pRUN`. Contains information necessary for communication. See `pMatlab_Init` for more details.

pMatlab_Init *Deprecated*

Called by `pRUN` to initialize pMatlab environment.

Syntax

`pMatlab_Init`

Description

Initializes variables required by the pMatlab library, such as number of processors, current processor's rank and which processor is the leader. All of the variables necessary for communication are stored in the `pMATLAB` structure.

Fields of the `pMATLAB` structure:

- `comm` - contains the MatlabMPI communicator
- `comm_size` - size of communicator, i.e. number of processors
- `my_rank` - rank of the local processor
- `leader` - indicates which rank is the leader, by default set to 0
- `pList` - list of ranks of participating processors

- `tag` - current message tag
- `tag_num` - number of messages sent; synchronized across all processors in `pList`

Fields to be potentially added in the future:

- `num_tasks` - number of tasks (scopes) created from the beginning of the program
- `curr_task` - current task (scope)
- `scopes` - contains a cell array of communication scopes; each entry is a struct with the current fields of the `pMATLAB` structure plus the `task_num` field

pMatlab_Finalize *Deprecated*

Called by `pRUN` to terminate pMatlab environment.

Syntax

```
pMatlab_Finalize
```

Description

Terminates pMatlab environment, i.e. exits non-leader MATLAB processes. This ensures that MATLAB processes are not orphaned on remote machines while leaving the leader process running.

pMatlab_ver

Display version number for pMatlab

Syntax

```
v = pMatlab_ver
```

Description

`v = pMatlab_ver` returns a string `v` containing the pMatlab version.

MPI_Abort

Aborts any currently running pMatlab or MatlabMPI program and blocks returning until all processes have ended. Automatically prepended to pMatlab scripts by `pRUN`.

Syntax

MPI_Abort

Description

Will abort any currently running pMatlab/MatlabMPI program by looking for leftover MATLAB processes and killing them. Cannot be used after `MatMPI_Delete_all`. Must be run in the directory from which the pMatlab/MatlabMPI programs was launched.

MatMPI_Delete_all

Deletes the MatMPI directory and its contents. Automatically prepended to pMatlab scripts by `pRUN`.

Syntax

```
MatMPI_Delete_all
```

MPI_Run *Deprecated*

Called by `pRUN` to launches a pMatlab v2.0 program. Used directly to run a pMatlab v1.0or MatlabMPI program.

Syntax

```
eval(MPI_Run(mfile, Ncpus, cpus))
```

Description

`mfile` is a string that contains the name of the pMatlab/MatlabMPI program to be launched, without the `.m` suffix.

`Ncpus` is an integer that specifies the number of processors to launch `mfile` onto.

`cpus` specifies what machines to launch `mfile` onto:

- `cpus = {};` Run all MATLAB processes on the local machine.
- `cpus = {'machine1' 'machine2' ...};` Specify names of machines on which to run. To run interactively, `machine1` must be the name of local machine.
- `cpus = {'machine1:dir1' 'machine2:dir2' ...};` Specify machines names and which directory to use for communication on each machine. Directories must be visible to both machines, i.e. crossmounted. Directories should be located on the local disk of their respective machines.
- `cpus = {'machine1:type' 'machine2:type'};` Specify machine names and the type of each machine. `type` can be either `'unix'` or `'pc'`. Default is `'unix'` (can be changed in `MatMPI_Comm_settings.m`)

`cpus = {'machine1:type:dir1' 'machine2:type:dir2'};` Specify machine names, communication directories, and the type of each machine.

SendMsg

Sends a set of variables from the current Pid to the Pid denoted by `dest`.

Syntax

```
SendMsg(dest,tag,var1,var2,...)
```

Description

`dest` is the numeric Pid of the destination.

`tag` is an integer or string used to differentiate multiple message being sent between the same Pids..

`var1,var2,...` are the variables to be sent.

RecvMsg

Receives a set of variables sent from the Pid denoted by `source`.

Syntax

```
[var1,var2,...] = RecvMsg(source,tag)
```

Description

`source` is the numeric Pid of the sender.

`tag` is an integer or string used to differentiate multiple message being sent between the same Pids..

`var1,var2,...` are the variables to be received.

Distributed matrices and matrix manipulation

Elementary distributed matrices

map/map

Map class constructor.


Syntax

```
p = map (GRID_SPEC, DIST_SPEC, PROC_LIST)
p = map (GRID_SPEC, DIST_SPEC, PROC_LIST, OVERLAP_SPEC)
```

Description

`map (GRID_SPEC, DIST_SPEC, PROC_LIST, OVERLAP_SPEC)` constructs a `map` object to be used as an input to a `dmat` constructor.

- **GRID_SPEC**: Vector of integers specifying how each dimension of a `dmat` is broken up. For example, if `GRID_SPEC = [2 3]`, the first dimension is broken up between 2 processors and the second dimension is broken up between 3 processors. The following figure illustrates how this grid example would break up a `dmat` given 6 processors using a block distribution.



The length of `GRID_SPEC` can be 2, 3, or 4 and must match the number of dimensions in the `dmat`.


- **DIST_SPEC**: Array of structures with two possible fields, `dist` and `b_size`, specifying the `dmat` distribution.
`DIST_SPEC.dist` is a string specifying the type of data distribution the `dmat` should use. Each entry in the array must have the `dist` field defined. The `dist` field can have three possible values:

- 'b': block
- 'c': cyclic
- 'bc': block-cyclic


Setting `DIST_SPEC` to `{ }` uses block distribution for all dimensions.

`DIST_SPEC.b_size` specifies the block size for block-cyclic distributions. If `DIST_SPEC.dist` is set to 'bc', then `DIST_SPEC.b_size` must also be defined. If `DIST_SPEC.dist` is set to 'b' or 'c', then `DIST_SPEC.b_size` does not have to be defined.


The following figure shows an example of the same `dmat` distributed over 4 processors using each of the three types of data distributions:



- `PROC_LIST`: Array of processor ranks specifying on which ranks the object should be distributed. Ranks are assigned column-wise (top-down, then left-right) to grid locations in sequential order.



- `OVERLAP_SPEC`: Optional. Vector of integers specifying amount of overlap between processors for each dimension. The following figure shows an example of a `dmat` distributed across four processors with 1 column of overlap between adjacent processors.



The length of `OVERLAP_SPEC` can be 2, 3, or 4 and must match the number of dimensions in the `dmat`. Only block distributions can have overlap.

`map` returns a data structure `p` which contains the following fields:

- `DIM`: the number of dimensions of the map (must equal the dimension of the `dmat`)
- `PROC_LIST`: the list of processor ranks on which the object should be distributed
- `DIST_SPEC`: the distribution specification for each dimension
- `GRID`: array of length `DIM` specifying how the object should be distributed

Examples

2D map, 2x2 grid, block-cyclic along rows and columns, block size 2 along rows, block size 3 along columns:

```
grid1 = [2 2];           % 2x2 grid
dist1(1).dist = 'bc';    % block-cyclic along dim 1 (rows)
dist1(1).b_size = 2;     % block size 2 along dim 1 (rows)
dist1(2).dist = 'bc';    % block-cyclic along dim 2 (columns)
dist1(2).b_size = 3;     % block size 3 along dim 2 (columns)
proc1 = [0:3];           % list of ranks 0 through 3
map1 = map(grid1, dist1, proc1);
```

2D map, 2x3 grid, cyclic along both rows and columns:

```
grid2 = [2 3];           % 2x3 grid
dist2(1).dist = 'c';     % cyclic along dim 1 (rows)
dist2(2).dist = 'c';     % cyclic along dim 2 (columns)
proc2 = [0:5];           % list of ranks 0 through 5
map2 = map(grid2, dist2, proc2);
```

2D map, 1x2 grid, block along rows, cyclic along columns:

```
grid3 = [1 2];           % 1x2 grid
dist3(1).dist = 'b';     % block along dim 1 (rows)
dist3(2).dist = 'c';     % cyclic along dim 2 (columns)
proc3 = [0:1];           % list of ranks 0 and 1
map3 = map(grid3, dist3, proc3);
```

3D map, 2x3x2 grid, block-cyclic along rows and columns with block size 2, cyclic along third dimension:

```
grid4 = [2 3 2];         % 2x3x2 grid
dist4(1).dist = 'bc';    % block-cyclic along dim 1 (rows)
dist4(1).b_size = 2;     % block size 2 along dim 1 (rows)
dist4(2).dist = 'bc';    % block-cyclic along dim 2 (columns)
dist4(2).b_size = 2;     % block size 2 along dim 2 (columns)
dist4(3).dist = 'c';     % cyclic along dim 3
proc4 = [0:11];          % list of ranks 0 through 12
map4 = map(grid4, dist4, proc4);
```

2D map, 1x4 grid, block along rows, cyclic along columns:

Distributed matrices and matrix manipulation

```
grid5 = [1 4];           % 1x4 grid
dist5(1).dist = 'b';      % block along dim 1 (rows)
dist5(2).dist = 'c';      % cyclic along dim 2 (columns)
proc5 = [0:3];           % list of ranks 0 through 3
map5 = map(grid5, dist5, proc5);
```

2D map, block along both dimensions, overlap in the column dimension of size 1 (1 column overlap):

```
grid6 = [2 2];           % 2x2 grid
dist6 = {};              % block along all dimensions
proc6 = [0 1];           % list of ranks 0 and 1
overlap6 = [0 1];        % overlap of 0 along dim 1 (rows)
                          % overlap of 1 along dim 2 (columns)
map6 = map(grid6, dist6, proc6, overlap6);
```

These examples show only how to create `map` objects. Refer to `dmat/ones`, `dmat/rand`, and `dmat/zeros` on how to create `dmat` objects using `map` objects.

map/zeros

Create a `dmat` of zeros.

Syntax

```
Y = zeros(N, P)
Y = zeros(M, N, P)
Y = zeros(M, N, Q, P)
Y = zeros(M, N, Q, R, P)
```

Description

`zeros(N, P)` returns an *N*-by-*N* `dmat` of zeros mapped according to the map specified by *P*.

`zeros(M, N, P)` returns an *M*-by-*N* `dmat` of zeros mapped according to the map specified by *P*.

`zeros(M, N, Q, P)` returns an *M*-by-*N*-by-*Q* `dmat` of zeros mapped according to the map specified by *P*.

`zeros(M, N, Q, R, P)` returns an *M*-by-*N*-by-*Q*-by-*R* `dmat` of zeros mapped according to the map specified by *P*.

Remarks

Dimension of the `dmat` must be consistent with the dimension of the map's grid.

map/ones

Create a `dmat` of all ones

Syntax

```
Y = ones (N, P)
Y = ones (M, N, P)
Y = ones (M, N, Q, P)
Y = ones (M, N, Q, R, P)
```

Description

`ones (N, P)` returns an *N*-by-*N* `dmat` of ones mapped according to the map specified by *P*.

`ones (M, N, P)` returns an *M*-by-*N* `dmat` of ones mapped according to the map specified by *P*.

`ones (M, N, Q, P)` returns an *M*-by-*N*-by-*Q* `dmat` of ones mapped according to the map specified by *P*.

`ones (M, N, Q, R, P)` returns an *M*-by-*N*-by-*Q*-by-*R* `dmat` of ones mapped according to the map specified by *P*.

Remarks

Dimension of the `dmat` must be consistent with the dimension of the map's grid.

map/rand

Create a `dmat` of uniformly distributed random numbers.

Syntax

```
Y = rand(N, P)
```

```
Y = rand(M, N, P)
```

```
Y = rand(M, N, Q, P)
```

```
Y = rand(M, N, Q, R, P)
```

Description

The `rand` function generates `dmats` of random numbers between 0 and 1 distributed uniformly.

`rand(N, P)` returns `N-by-N` `dmat` of random numbers mapped according to the map specified by `P`.

`rand(M, N, P)` returns `M-by-N` `dmat` of random numbers mapped according to the map specified by `P`.

`rand(M, N, Q, P)` returns an `M-by-N-by-Q` `dmat` of random numbers mapped according to the map specified by `P`.

`rand(M, N, Q, R, P)` returns an `M-by-N-by-Q-by-R` `dmat` of random numbers mapped according to the map specified by `P`.

Remarks

Dimension of the `dmat` must be consistent with the dimension of the map's grid.

Calls the MATLAB `rand` function to create each local part of the `dmat`. Thus, the resulting array will **not** be the same as a `double` random array of the same dimensions.

Basic array information

dmat/size

Size of the dmat.

Syntax

```
d = size(X)
[m, n] = size(X)
[d1, d2, d3, ..., dn] = size(X)
```

Description

`d = size(X)` returns the size of each dimension of dmat X in vector d.

`[m,n] = size(X)` returns the size of dmat X in separate variables m and n.

`[d1,d2,d3,...,dn] = size(X)` returns the sizes of each dimension of x in separate variables.

Remarks

If `A = zeros(m, n, q, p1)` and `B = zeros(m, n, q, p2)`, where p1 and p2 are different maps, `size(A)` and `size(B)` return the same results.

dmat/ndims

Number of dimension of the dmat.

Syntax

```
n = ndims(A)
```

Description

`n = ndims(A)` returns the number of dimensions in the dmat A. The number of dimensions in a dmat is always greater than or equal to 2.

Remarks

`ndims(A)` is `length(size(A))`.

dmat/display

Display `dmat`.

Syntax

```
display(D)
```

Description

`display(D)` aggregates the `D` onto the leader process and displays the entire contents of `D` on the leader process. On remote processes, `display(D)` displays only the local portion of `D`.

`display(D)` is also called for `D` when a semicolon is not used to terminate a statement.

Remarks

Note that `display` incurs communication overhead to aggregate `D` onto the leader processor.

map/display

Display `map` object.

Syntax

```
display(M)
```

Description

`display(M)` displays the contents of the `map` object.

Remarks

`display(M)` is also called for `M` when a semicolon is not used to terminate a statement.

Distributed array information

dmat/global_block_range

Returns the ranges of global indices local to the current processor for a given `dmat`.

Syntax

```
I = global_block_range(D, DIM)
[I1, I2, ..., IN] = global_block_range(D)
```

Description

`I = global_block_range(D, DIM)` Returns the global index range of the `dmat` `D` local to the current processor in the specified dimension, `DIM`.

`[I1, I2, ..., IN] = global_block_range(D)` Returns the global index range of the `dmat` `D` local to the current processor for all `N` dimensions of `D`.

The global index range for each dimension is returned as a 2-element vector. The first element in the vector represents the starting global index and the second element represents the ending index.

Examples

Let `Ncpus` be 4:

```
P = map([1 Ncpus], {}, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_block_range(D);
```

For each rank, `I1` contains:

Rank	I1 (1)	I1 (2)
0	1	100
1	1	100
2	1	100
3	1	100

For each rank, `I2` contains:

Rank	I2 (1)	I2 (2)
0	1	25
1	26	50
2	51	75
3	76	100

dmatrix/global_block_ranges

Returns the ranges of global indices for all processors in the map of `dmatrix D`.

Syntax

```
I = global_block_ranges(D, DIM)
[I1, I2, ..., IN] = global_block_ranges(D)
```

Description

`I = global_block_ranges(D, DIM)` Returns the global index ranges of the `dmatrix D` for all processors in the specified dimension, `DIM`.

`[I1, I2, ..., IN] = global_block_ranges(D)` Returns the global index range of the `dmatrix D` for all processors in all dimensions of `D`.

For each dimension, the indices are returned as a matrix `I` of size `NUM_PROCS_IN_GRIDx3`. Each line of the returned matrix, `I(i, :)` contains the following information:

```
[PROCESSOR_RANK    START_INDEX    END_INDEX]
```

Examples

Let `Ncpus` be 4:

```
P = map([1 Ncpus], {}, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_block_ranges(D);
```

On every rank, `I1` contains:

<code>I1(1)</code>	<code>I1(2)</code>	<code>I1(3)</code>
0	1	100
2	1	100
1	1	100
3	1	100

On every rank, `I2` contains:

<code>I2(1)</code>	<code>I2(2)</code>	<code>I2(3)</code>
0	1	25
2	26	50
1	51	75
3	76	100

Remarks

The difference between `global_block_range` and `global_block_ranges` is subtle, but important. `global_block_range` returns a single vector containing the index range for *only* that particular processor. `global_block_ranges` returns a matrix that contains the index ranges for *every* processor.

dmatrix/global_ind

Returns the global indices local to the current processor.

Syntax

```
I = global_ind(D, DIM)
[I1, I2, ..., IN] = global_ind(D)
```

Description

`I = global_ind(D, DIM)` Returns the global indices of the `dmatrix D` local to the current processor in the specified dimension, `DIM`.

`[I1, I2, ..., IN] = global_ind(D)` Returns the global indices of the `dmatrix D` local to the current processor in all dimensions of `D`.

The global indices for each dimension are returned as a vector.

Examples

Let `Ncpus` be 4:

```
P = map([1 Ncpus], {}, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_ind(D);
```

For each rank, `I1` contains:

Rank	<code>I1(:)</code>
0	1 2 3 ... 99 100
1	1 2 3 ... 99 100
2	1 2 3 ... 99 100
3	1 2 3 ... 99 100

For each rank, `I2` contains:

Rank	<code>I2(:)</code>
0	1 2 3 ... 24 25
1	26 27 28 ... 49 50
2	51 52 53 ... 74 75
3	76 77 78 ... 99 100

dmatrix/global_inds

Returns the global indices for all processors in the map of `dmatrix D`.

Syntax

```
I = global_inds(D, DIM)
[I1, I2, ..., IN] = global_inds(D)
```

Description

`global_inds(D, DIM)` Returns global indices of the `dmatrix D` for all processors in the specified dimension, `DIM`.

`global_inds(D)` Returns global indices of the `dmatrix D` for all processors in all dimensions of `D`.

For each dimension, the indices are returned as a matrix `I` of size `NUM_PROCS_IN_GRID X MAX_LOCAL_INDICES`. Each line of the returned matrix `I`, `I(i, :)`, contains the following information:

```
[PROCESSOR_RANK  IND1  IND2  ...  INDn]
```

To ensure that all rows in the return index are the same, the indices matrix is appended with extra zeros where there are not enough indices.

Examples

Let `Ncpus` be 4:

```
P = map([1 Ncpus], {}, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_ind(D);
```

On every rank, `I1` contains:

<code>I1(1)</code>	<code>I1(2:end)</code>
0	1 2 3 ... 99 100
1	1 2 3 ... 99 100
2	1 2 3 ... 99 100
3	1 2 3 ... 99 100

On every rank, `I2` contains:

<code>I2(1)</code>	<code>I2(2:end)</code>
0	1 2 3 ... 24 25
1	26 27 28 ... 49 50
2	51 52 53 ... 74 75
3	76 77 78 ... 99 100

Remarks

The difference between `global_ind` and `global_inds` is subtle, but important. `global_ind` returns a single vector containing the indices for only that particular processor. `global_inds` returns a matrix that contains the indices for every processor.

dmatrix/global_range

Returns the ranges of global indices `dmatrix D` of local to the current processor. Returns the same range as `global_block_range` if `D` is block distributed, returns subranges for block-cyclic and cyclic distributions.

Syntax

```
I = global_range(D, DIM)
[I1, I2, ..., IN] = global_range(D)
```

Description

`I = global_range(D, DIM)` Returns the global index range of the `dmatrix D` local to the current processor in the specified dimension, `DIM`.

`[I1, I2, ..., IN] = global_range(D)` Returns the global index range of the `dmatrix D` local to the current processor in all dimensions of `D`.

For each dimension, the indices are returned as a matrix `I`. Each line of the returned matrix, `I(i, :)`, contains the following information:

```
[START_INDEX_1 END_INDEX_1 START_INDEX_2 END_INDEX_2 ...]
```

Examples

Let `Ncpus` be 4:

```
dist(1).dist = 'b';
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_range(D);
```

For each rank, `I1` contains:

Rank	I1
0	[1 50]
1	[51 100]
2	[1 50]
3	[51 100]

For each rank, `I2` contains:

Rank	I2
0	[1 50]
1	[1 50]
2	[51 100]
3	[51 100]

Let Ncpus be 4:

```
dist(1).dist = 'c';
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_range(D);
```

For each rank, I1 contains:

Rank	I1
0	[1 1 3 3 5 5 ... 97 97 99 99]
1	[2 2 4 4 6 6 ... 98 98 100 100]
2	[1 1 3 3 5 5 ... 97 97 99 99]
3	[2 2 4 4 6 6 ... 98 98 100 100]

For each rank, I2 contains:

Rank	I2
0	[1 50]
1	[1 50]
2	[51 100]
3	[51 100]

Let Ncpus be 4:

```
dist(1).dist = 'bc';
dist(1).b_size = 4;
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_range(D);
```

For each rank, I1 contains:

Rank	I1
0	[1 4 9 12 ... 89 92 97 100]
1	[5 8 13 16 ... 93 96 0 0]
2	[1 4 9 12 ... 89 92 97 100]
3	[5 8 13 16 ... 93 96 0 0]

For each rank, I2 contains:

Rank	I2
0	[1 50]
1	[1 50]
2	[51 100]
3	[51 100]

dmat/global_ranges

Returns the ranges of global indices for all processors in the map of dmat *D*. Returns the same range as `global_block_ranges` if *D* is block distributed, returns subranges for block-cyclic and cyclic distributions.

Syntax

```
I = global_ranges(D, DIM)
[I1, I2, ..., IN] = global_ranges(D)
```

Description

`I = global_ranges(D, DIM)` Returns the global index ranges of the dmat *D* for all processors in the specified dimension, *DIM*.

`[I1, I2, ..., IN] = global_ranges(D)` Returns the global index range of the dmat *D* for all processors in all dimensions of *D*.

For each dimension, the indices are returned as a matrix *I* of size

`NUM_PROCS_IN_GRID × NUM_BLOCK_BOUNDARIES`. Each line of the returned matrix, `I(i, :)`, contains the following information:

```
[PROCESSOR_RANK START_INDEX_1 END_INDEX_1 START_INDEX_2 END_INDEX_2 ...]
```

Examples

Let *Ncpus* be 4:

```
dist(1).dist = 'b';
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_ranges(D);
```

On every rank, *I1* contains:

<i>I1</i> (1)	<i>I1</i> (2:end)
0	1 50
2	1 50
1	51 100
3	51 100

On every rank, *I2* contains:

<i>I2</i> (1)	<i>I2</i> (2:end)
0	1 50
2	51 100
1	1 50
3	51 100

Let Ncpus be 4:

```
dist(1).dist = 'c';
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_ranges(D);
```

On every rank, I1 contains:

I1 (1)	I1 (2:end)
0	1 1 3 3 5 5 ... 97 97 99 99
2	1 1 3 3 5 5 ... 97 97 99 99
1	2 2 4 4 6 6 ... 98 98 100 100
3	2 2 4 4 6 6 ... 98 98 100 100

On every rank, I2 contains:

I2 (1)	I2 (2:end)
0	1 50
2	51 100
1	1 50
3	51 100

Let Ncpus be 4:

```
dist(1).dist = 'bc';
dist(1).b_size = 4;
dist(2).dist = 'b';
P = map([Ncpus/2 Ncpus/2], dist, 0:Ncpus-1);
D = zeros(100, 100, P);
[I1, I2] = global_ranges(D);
```

On every rank, I1 contains:

I1 (1)	I1 (2:end)
0	1 4 9 12 ... 89 92 97 100
2	1 4 9 12 ... 89 92 97 100
1	5 8 13 16 ... 93 96 0 0
3	5 8 13 16 ... 93 96 0 0

On every rank, I2 contains:

I2 (1)	I2 (2:end)
0	1 50
2	51 100
1	1 50
3	51 100

Remarks

If processors in the same dimension have different number of blocks, the block boundaries are padded with zeros for the processors that have fewer blocks.

map/inmap

Checks if a processor is in the map.

Syntax

```
b = inmap(m, r)
```

Description

`b = inmap(m, r)` checks if processor rank `r` is in map `m`. Returns TRUE for Boolean if rank `r` is in the map, FALSE otherwise.

Matrix manipulation

dmat/find

Find indices of nonzero elements in a `dmat`.

Syntax

```
[I, J] = find(X)
```

Description

`[I, J] = find(X)` returns the row and column indices of nonzero elements of the `dmat` `X`.

Remarks

Currently supports only `[I, J] = find(X)` calling convention. Only works on 2D `dmats`.

`find` requires every processor to send its results to every other processor, thus can incur a significant amount of communication overhead.

Distributed matrix manipulation

dmat/agg

Aggregates the parts of a `dmat` on the leader processor.

Syntax
$$A = \text{agg}(D)$$
Description

`A = agg(D)` aggregates the parts of a `dmat` `D` into a whole and returns it as a regular double matrix, `A`. If the current processor is the leader, returns the aggregated matrix. Otherwise, returns the local part of `D`.

Remarks

Currently, it doesn't matter if the leader is in the map – the global matrix is returned on the leader, regardless.

Note that `agg` incurs communication overhead to aggregate `D` onto the leader processor.

Since `A` on the leader processor contains the entire contents of `D` but on all other processors contains only the local portion of `D`, `A` will have different values and sizes on each processor. Thus, `agg` should be used with caution.

dmatrix/agg_all

Aggregates the parts of a `dmatrix` onto all processors.

Syntax

```
A = agg_all(D)
```

Description

`A = agg_all(D)` aggregates the parts of a `dmatrix` `D` onto all processors in the map of `D` and returns a regular double matrix `A`.

Remarks

Unlike `agg`, `agg_all` creates a result that is consistent in size and values across all processors. However, because `agg_all` causes all processors to communicate with all processors, `agg_all` can incur a significant amount of communication and should be used with caution.

dmatrix/local

Returns the local part of the `dmatrix`.

Syntax


```
D_local = local(D)
```

Description

`D_local = local(D)` Returns the local part of the `dmatrix` `D` on the current processor.

Examples

The following diagram shows four processors obtaining their respective local parts of the `dmatrix`, `D`, and copying the contents into a local variable, `D_local`. Note that `D_local` exists on each processor but contains different data.



dmat/put_local

Assigns new data to the local part of the dmat.

Syntax


```
D = put_local(D, D_LOCAL)
```

Description

$D = put_local(D, D_LOCAL)$ assigns D_LOCAL to the local part of the dmat D .

Examples

The following diagram shows four processors each writing a local matrix, D_local , into their respective parts of a dmat, D . Note that D_local on each processor must be the same size as the local portion of their respective parts of D .



dmat/synch

Synchronize the overlapped data in a `dmat`.

Syntax


```
D = synch(D)
```

Description

`D = synch(D)` If overlap is present, the owner processor of the overlapping data sends its data to the processor that has a copy of the overlapping data. No-op if there is no overlap.

Remarks

The owner is the processor with the higher index in the grid in the corresponding dimension. For example, if the overlap is in the second dimension the owner is the processor in the column of the grid with the higher index.

***dmat/sync2***

Perform an operation on the overlapped data with the result that both processors hold the final data locally.

Syntax


```
D = sync2(D)
```

Description

`D = sync2(D)` If overlap is present, each processor sends its overlapping data to its neighbor. Currently only the addition operator is supported, so each processor adds the received data to its given data, the result being that all processors have the sum of their overlap. No-op if there is no overlap.

Remarks

The owner is the processor with the higher index in the grid in the corresponding dimension. For example, if the overlap is in the second dimension the owner is the processor in the column of the grid with the higher index.

***remap***

Remaps a `dmat` with a new map.

Syntax

```
remap(X, NEW_MAP)
```

Description

`remap(X, NEW_MAP)` takes a `dmat` `X` and redistributes it according to the specified map `NEW_MAP`.

dmat/subsasgn

Subscripted assignment to a distributed object. Overloaded method for $A(I) = B$. Should not be called directly.

Syntax

```
A = subsasgn(A, S, B)
```

Description

`A = subsasgn(A, S, B)` Subscripted assignment of `B` (right hand side) to `A` (left hand side). `A` is of type `dmat`. `B` can be of type `dmat` or `double`. `S` is a structure array with the fields:

- `type`: String containing '()', '{}', or '.' specifying the subscript type. Currently only supports '()'.
- `subs`: Cell array or string containing the actual subscripts.

Remarks

In cases where $A(I)$ and/or B are distributed across multiple processors, `subsasgn` will automatically transfer the appropriate data between processors.

dmatrix/subsref

Subscripted reference. Overloaded method for $A(I)$. Should not be called directly.

Syntax

```
B = subsref(A, S)
```

Description

$B = \text{subsref}(A, S)$ Subscripted reference on a $\text{dmatrix } A$. S is a structure array with the fields:

- `type`: String containing '()', '{}', or '.' specifying the subscript type.
- `subs`: Cell array or string containing the actual subscripts.

Remarks:

Currently, `subsref` will only return a standalone dmatrix in the following cases:

- $A(:, :)$ – Refers to the entire contents of the $\text{dmatrix } A$
- $A(i, j)$ – Refers to a single element in the $\text{dmatrix } A$. Returns a new dmatrix containing the value at (i, j) stored on the processor which contained that element in A .

In all other cases, `subsref` will not produce a “standalone” dmatrix , i.e. the resulting dmatrix can not be directly used as an input to any pMatlab function, with the exception of `local`.

map/subsasgn

Subscripted assignment. Should not be called directly.

Syntax

$$A = \text{subsasgn}(A, S, B)$$
Description

$A = \text{subsasgn}(A, S, B)$ Subscripted assignment of B (right hand side) to A (left hand side).

$A.\text{FIELD} = B$ allows the fields of a `map` object A to be assigned using the `'.'` notation (complies with structure behavior).

Remarks

This functionality might be deprecated from the final API, to limit control the user has of private members of the MAP object.

map/subsref

Subscripted reference. Should not be called directly.

Syntax

$$B = \text{subsref}(A, S)$$
Description

$A = \text{subsref}(A, S, B)$ Subscripted assignment of A (right hand side) to B (left hand side).

$B = A.\text{FIELD}$ allows the fields of a `map` object A to be referenced using the `'.'` notation (complies with structure behavior).

Remarks

This functionality might be deprecated from the final API, to limit control the user has of private members of the MAP object. `subsref` might be replaced by `get` functions.

transpose_grid

Redistributes a `dmatrix` by transposing its grid.

Syntax

```
A = transpose_grid(B)
```

Description

`A = transpose_grid(B)` creates a `dmatrix` `A` that has the same contents as the `dmatrix` `B`, except that the grid for `A`'s map is the transpose of the grid for `B`'s map. The contents of `B` are automatically redistributed to `A`. `transpose_grid` is only supported for 2D `dmatrix`s.

`transpose_grid` is optimized to redistribute row-distributed `dmatrix`s into column-distributions and vice versa. For example, suppose `B`'s map had a grid specification of `[1 4]`. Then `A`'s map will have a grid specification of `[4 1]`. Note that `B` must be block distributed in both dimensions with no overlap.

For all other distributions, e.g. `dmatrix`s with row and column distributions, overlap, etc., `subsasgn` will be used to redistribute `B`.

Elementary math functions

Trigonometric

dmat/sin, dmat/cos, dmat/tan, dmat/sec, dmat/csc, dmat/cot

Performs the trigonometric operation in radians.

Syntax

`Y = sin(X)`

`Y = cos(X)`

`Y = tan(X)`

`Y = sec(X)`

`Y = csc(X)`

`Y = cot(X)`

Description

`Y = sin(X)` is the sine of the elements of the distributed array `x`, expressed in radians.

`Y = cos(X)` is the cosine of the elements of the distributed array `x`, expressed in radians.

`Y = tan(X)` is the tangent of the elements of the distributed array `x`, expressed in radians.

`Y = sec(X)` is the secant of the elements of the distributed array `x`, expressed in radians.

`Y = csc(X)` is the cosecant of the elements of the distributed array `x`, expressed in radians.

`Y = cot(X)` is the cotangent of the elements of the distributed array `x`, expressed in radians.

dmatrix/sind, dmatrix/cosd, dmatrix/tand, dmatrix/secd, dmatrix/cscd, dmatrix/cotd

Performs the trigonometric operation in degrees.

Syntax

`Y = sind(X)`

`Y = cosd(X)`

`Y = tand(X)`

`Y = secd(X)`

`Y = cscd(X)`

`Y = cotd(X)`

Description

`Y = sind(X)` is the sine of the elements of the distributed array `x`, expressed in degrees.

`Y = cosd(X)` is the cosine of the elements of the distributed array `x`, expressed in degrees.

`Y = tand(X)` is the tangent of the elements of the distributed array `x`, expressed in degrees.

`Y = secd(X)` is the secant of the elements of the distributed array `x`, expressed in degrees.

`Y = cscd(X)` is the cosecant of the elements of the distributed array `x`, expressed in degrees.

`Y = cotd(X)` is the cotangent of the elements of the distributed array `x`, expressed in degrees.

dmatrix/sinh, dmatrix/cosh, dmatrix/tanh, dmatrix/sech, dmatrix/csch, dmatrix/coth

Performs the hyperbolic trigonometric operation.

Syntax

`Y = sinh(X)`

`Y = cosh(X)`

`Y = tanh(X)`

`Y = sech(X)`

`Y = csch(X)`

`Y = coth(X)`

Description

`Y = sinh(X)` is the hyperbolic sine of the elements of the distributed array `x`.

`Y = cosh(X)` is the hyperbolic cosine of the elements of the distributed array `x`.

`Y = tanh(X)` is the hyperbolic tangent of the elements of the distributed array `x`.

`Y = sech(X)` is the hyperbolic secant of the elements of the distributed array `x`.

`Y = csch(X)` is the hyperbolic cosecant of the elements of the distributed array `x`.

`Y = coth(X)` is the hyperbolic cotangent of the elements of the distributed array `x`.

dmat/asin, dmat/acos, dmat/atan, dmat/asec, dmat/acsc, dmat/acot

Performs the inverse trigonometric operation in radians.

Syntax

$Y = \text{asin}(X)$

$Y = \text{acos}(X)$

$Y = \text{atan}(X)$

$Y = \text{asec}(X)$

$Y = \text{acsc}(X)$

$Y = \text{acot}(X)$

Description

$Y = \text{asin}(X)$ is the inverse sine of the elements of the distributed array x , expressed in radians. Complex results are obtained if $\text{abs}(x) > 1.0$ for some element.

$Y = \text{acos}(X)$ is the inverse cosine of the elements of the distributed array x , expressed in radians. Complex results are obtained if $\text{abs}(x) > 1.0$ for some element.

$Y = \text{atan}(X)$ is the inverse tangent of the elements of the distributed array x , expressed in radians.

$Y = \text{asec}(X)$ is the inverse secant of the elements of the distributed array x , expressed in radians.

$Y = \text{acsc}(X)$ is the inverse cosecant of the elements of the distributed array x , expressed in radians.

$Y = \text{acot}(X)$ is the inverse cotangent of the elements of the distributed array x , expressed in radians.

dmatrix/asind, dmatrix/acosd, dmatrix/atand, dmatrix/asecd, dmatrix/acscd, dmatrix/acotd

Performs the inverse trigonometric operation in degrees.

Syntax

```
Y = asind(X)
Y = acosd(X)
Y = atand(X)
Y = asecd(X)
Y = acscd(X)
Y = acotd(X)
```

Description

$Y = \text{asin}(X)$ is the inverse sine of the elements of the distributed array x , expressed in degrees.

$Y = \text{acos}(X)$ is the inverse cosine of the elements of the distributed array x , expressed in degrees.

$Y = \text{atan}(X)$ is the inverse tangent of the elements of the distributed array x , expressed in degrees.

$Y = \text{asec}(X)$ is the inverse secant of the elements of the distributed array x , expressed in degrees.

$Y = \text{acsc}(X)$ is the inverse cosecant of the elements of the distributed array x , expressed in degrees.

$Y = \text{acot}(X)$ is the inverse cotangent of the elements of the distributed array x , expressed in degrees.

dmatrix/asinh, dmatrix/acosh, dmatrix/atanh, dmatrix/asech, dmatrix/acsch, dmatrix/acoth

Performs the inverse hyperbolic trigonometric operation.

Syntax

`Y = asinh(X)`

`Y = acosh(X)`

`Y = atanh(X)`

`Y = asech(X)`

`Y = acsch(X)`

`Y = acoth(X)`

Description

`Y = asinh(X)` is the inverse hyperbolic sine of the elements of the distributed array `x`.

`Y = acosh(X)` is the inverse hyperbolic cosine of the elements of the distributed array `x`.

`Y = atanh(X)` is the inverse hyperbolic tangent of the elements of the distributed array `x`.

`Y = asech(X)` is the inverse hyperbolic secant of the elements of the distributed array `x`.

`Y = acsch(X)` is the inverse hyperbolic cosecant of the elements of the distributed array `x`.

`Y = acoth(X)` is the inverse hyperbolic cotangent of the elements of the distributed array `x`.

Exponential

dmatrix/exp

Exponential.

Syntax

$$Y = \exp(X)$$

Description

$\exp(X)$ is the exponential of the elements of x , e to the x .

For complex $Z=X+i*Y$, $\exp(Z) = \exp(X) * (\cos(Y) + i*\sin(Y))$.

dmatrix/expm1

Compute $\exp(x) - 1$ accurately.

Syntax

$$Y = \expm1(X)$$

Description

$\expm1(X)$ computes $\exp(X) - 1$, compensating for the roundoff in $\exp(X)$ for distributed array x . For small x , $\expm1(X)$ is approximately x , whereas $\exp(X) - 1$ can be zero.

dmatrix/log

Natural logarithm.

Syntax

$$Y = \log(X)$$

Description

$\log(X)$ is the natural logarithm of the elements of distributed array x . Complex results are produced if x is not positive.

dmatrix/log1p

Compute $\log(1+x)$ accurately.

Syntax

$$Y = \log1p(X)$$

Description

$\log1p(X)$ computes $\log(1+x)$, compensating for the roundoff in $1+x$ for distributed array x . For small x , $\log1p(X)$ is approximately x , whereas $\log(1+x)$ can be zero.

dmatrix/log10

Compute common (base 10) logarithm.

Syntax

$$Y = \log10(X)$$

Description

$\log10(X)$ is the base 10 logarithm of the elements of distributed array x . Complex results are produced if x is not positive.

dmatrix/log2

Compute base 2 logarithm.

Syntax

$$Y = \log2(X)$$

Description

$\log2(X)$ is the base 2 logarithm of the elements of distributed array x .

dmatrix/pow2

Compute base 2 power.

Syntax

$$Y = \text{pow2}(X)$$
Description

$\text{pow2}(X)$ for each element of distributed array x is 2 raised to the power x .

dmatrix/realpow

Element by element power that will error out on complex result.

Syntax

$$R = \text{realpow}(P, Q)$$
Description

$\text{realpow}(P, Q)$ denotes element-by-element powers. The array dimensions must be the same, unless the non-distributed argument is a scalar. If both P and Q are *dmatrix*s, their maps must be the same. The *dmatrix* input must not be the only scalar, since operations with a distributed scalar incur significant communication overhead.

dmatrix/reallog

Natural log of real number.

Syntax

$$Y = \text{reallog}(X)$$
Description

$\text{reallog}(X)$ is the natural logarithm of the elements of distributed array x . An error is produced if x is not positive.

dmatrix/real_sqrt

Square root of number greater than or equal to zero.

Syntax

```
Y = real_sqrt(X)
```

Description

`real_sqrt(X)` is the square root of the elements of distributed array `x`. An error is produced if `x` is not positive.

dmatrix/sqrt

Square root.

Syntax

```
Y = sqrt(X)
```

Description

`sqrt(X)` is the square root of the elements of distributed array `x`. Complex results are produced if `x` is not positive.

Complex

dmatrix/abs

Absolute value.

Syntax

$$Y = \text{abs}(X)$$

Description

$\text{abs}(X)$ is the absolute values of the elements of distributed array x .

dmatrix/angle

Phase angle

Syntax

$$Y = \text{angle}(X)$$

Description

$\text{angle}(X)$ returns the phase angles, in radians, of a distributed array x with complex elements.

dmatrix/complex

Construct complex dmatrix from real dmatrix .

Syntax

$$C = \text{complex}(A)$$

$$C = \text{complex}(A, B)$$

Description

$C = \text{complex}(A)$ for real A returns the complex dmatrix C with real part A and all zero imaginary part.

$C = \text{complex}(A, B)$ returns the complex dmatrix $A + Bi$. A and B must have the same mapping.

dmatrix/conj

Complex conjugate.

Syntax
$$Y = \text{conj}(X)$$
Description

`conj(X)` is the complex conjugate of distributed array `x`.

For a complex `x`, $\text{CONJ}(X) = \text{REAL}(X) - i * \text{IMAG}(X)$.

dmatrix/imag

Complex imaginary part.

Syntax
$$Y = \text{imag}(X)$$
Description

`imag(X)` is the imaginary part of distributed array `x`.

dmatrix/real

Complex real part.

Syntax
$$Y = \text{real}(X)$$
Description

`real(X)` is the real part of distributed array `x`.

Rounding and remainder

dmatrix/fix

Round towards zero.

Syntax

$$Y = \text{fix}(X)$$

Description

`fix(X)` rounds the elements of distributed array `x` to the nearest integers towards zero.

dmatrix/floor

Round towards minus infinity.

Syntax

$$Y = \text{floor}(X)$$

Description

`floor(X)` rounds the elements of the distributed array `x` to the nearest integers towards minus infinity.

dmatrix/ceil

Round towards plus infinity.

Syntax

$$Y = \text{ceil}(X)$$

Description

`ceil(X)` rounds the elements of distributed array `x` to the nearest integers towards infinity.

dmatrix/round

Round towards nearest integer.

Syntax

`Y = round(X)`

Description

`round(X)` rounds the elements of distributed array `x` to the nearest integers.

dmatrix/sign

Signum function.

Syntax

`Y = sign(X)`

Description

`sign(X)` For each element of distributed array `X`, `sign(X)` returns 1 if the element is greater than zero, 0 if it equals zero and -1 if it is less than zero.

For the nonzero elements of complex `X`, `sign(X) = X ./ ABS(X)`.

Operators and special characters

dmatrix/plus

+ Plus.

Syntax

$R = P + Q$

$R = \text{plus}(P, Q)$

Description

$R = P + Q$ or $R = \text{plus}(P, Q)$ performs an element-wise addition between P and Q . If both P and Q are `dmatrix`s, then their dimensions and maps must be equal. If only one of P and Q is a `dmatrix`, then their dimensions must be the same or the non-`dmatrix` input must be a scalar. `dmatrix` inputs can only be scalar if both inputs are `dmatrix`s and have the same map.

dmatrix/minus

- Minus.

Syntax

$R = P - Q$

$R = \text{minus}(P, Q)$

Description

$R = P - Q$ or $R = \text{minus}(P, Q)$ performs an element-wise subtraction between P and Q (subtract Q from P). If both P and Q are `dmatrix`s, then their dimensions and maps must be equal. If only one of P and Q is a `dmatrix`, then their dimensions must be the same or the non-`dmatrix` input must be a scalar. `dmatrix` inputs can only be scalar if both inputs are `dmatrix`s and have the same map.

dmat/mtimes

* Matrix multiply.

Syntax

```
C = A * B
C = mtimes(A, B)
```

Description

`C = A * B` or `C = mtimes(A, B)` multiplies two matrices together. `Mtimes` is an inner product matrix multiply and assumes that `A` is row mapped and `b` is column mapped. If `B` is row mapped, it is remapped in order to perform the matrix multiplication. For `A` row mapped, the function accepts `B` matrices that are double (not mapped), column mapped or row mapped. For a matrix `A` that is not mapped (doubles), `B` must be column mapped.

Remarks

Overlaps have not been tested.

dmat/times

.* Array multiply.

Syntax

```
C = A .* B
C = times(A, B)
```

Description

`C = A .* B` or `C = times(A, B)` performs an element-wise multiplication between `A` and `B`. If both `A` and `B` are `dmat`s, then their dimensions and maps must be equal. If only one of `A` and `B` is a `dmat`, then their dimensions must be the same or the non-`dmat` input must be a scalar. `dmat` inputs can only be scalar if both inputs are `dmat`s and have the same map.

dmatrix/power

`.^` Array power.

Syntax

```
R = P .^ Q
R = power(P, Q)
```

Description

`R = P .^ Q` or `R = power(P, Q)` performs an element-wise power operation between `P` and `Q`. If both `P` and `Q` are `dmatrix`s, then their dimensions and maps must be equal. If only one of `P` and `Q` is a `dmatrix`, then their dimensions must be the same or the non-`dmatrix` input must be a scalar. `dmatrix` inputs can only be scalar if both inputs are `dmatrix`s and have the same map.

dmatrix/ldivide

`.\` Left array divide.

Syntax

```
R = P .\ Q
R = ldivide(P, Q)
```

Description

`R = P .\ Q` or `R = ldivide(P, Q)` performs an element-wise left divide between `P` and `Q` (divide `Q` by `P`). If both `P` and `Q` are `dmatrix`s, then their dimensions and maps must be equal. If only one of `P` and `Q` is a `dmatrix`, then their dimensions must be the same or the non-`dmatrix` input must be a scalar. `dmatrix` inputs can only be scalar if both inputs are `dmatrix`s and have the same map.

dmat/rdivide

`./` Right array divide.

Syntax

```
R = P ./ Q
R = rdivide(P, Q)
```

Description

`R = P ./ Q` or `R = rdivide(P, Q)` performs an element-wise right divide between `P` and `Q` (divide `P` by `Q`). If both `P` and `Q` are `dmats`, then their dimensions and maps must be equal. If only one of `P` and `Q` is a `dmat`, then their dimensions must be the same or the non-`dmat` input must be a scalar. `dmat` inputs can only be scalar if both inputs are `dmats` and have the same map.

dmat/eq

`==` Equal.

Syntax

```
A == B
```

Description

`A == B` compares the dimensions, maps, sizes and data of `A` and `B`.

If `A` and `B`'s maps, dimensions, and sizes agree then the output is a `dmat` with 0 where elements are not equal and 1 where elements are equal (similar to serial MATLAB).

If the maps are not equal, then a 0 is returned regardless of any other properties.

If the maps agree but the dimensions and/or sizes are not equal, then an error is thrown (analogous to serial MATLAB behavior).

map/eq

`==` Equal.

Syntax

$$A == B$$
Description

`A == B` compares member variables of maps `A` and `B`. If all are the same then `TRUE` is returned, otherwise `FALSE` is returned.

dmat/gt

`>` Greater than.

Syntax

$$A > B$$
Description

`A > B` compares each element of `dmat A` to scalar `B`. Returns a `dmat` with each entry equal to 0 if original entry was `< B`, and 1 otherwise. Calls the MATLAB `gt` on the local part of `A`.

map/ne

`~=` Not equal.

Syntax

$$A ~= B$$
Description

`A ~= B` Returns `FALSE` if two maps are equal, `TRUE` otherwise. Two maps are equal if their grids are equal.

dmatrix/transpose**.'** Transpose**Syntax**

```
A = B.'
```

```
A = transpose(B)
```

Description

B.' is the non-conjugate transpose of a distributed matrix. The input distributed array must be a matrix. The distribution of the distributed array is limited to block in 1, 2nd, or both dimensions.

dmatrix/transpose**'** Complex conjugate transpose**Syntax**

```
A = B'
```

```
A = ctranspose(B)
```

Description

B.' is the complex conjugate transpose of a distributed matrix. The input distributed array must be a matrix. The distribution of the distributed array is limited to block in 1st, 2nd, or both dimensions.

Sparse matrices

map/sparse

Create a sparse dmat.

Syntax

```
S = sparse([I], [J], [S], M, N, NZMAX, P)
S = sparse([I], [J], [S], M, N, P)
S = sparse([I], [J], [S], P)
S = sparse(M, N, P)
```

Description

`S = sparse(I, J, S, M, N, NZMAX, P)` generates an M-by-N sparse dmat distributed according to map P with space allocated for NZMAX nonzeros (note that NZMAX applies to the overall dmat, not to individual processors. NZMAX will be distributed as evenly as possible over all processors).

The rows of `[I, J, S]` are intended to be used to initialize the non-zero values of the matrix. However, `sparse` currently does not support the use of `I, J,` and `S` in pMatlab; they are kept to remain consistent with the `sparse` function built into MATLAB.

There are four ways that `sparse` can be called:

- `S = sparse([I], [J], [S], M, N, NZMAX, P)`
- `S = sparse([I], [J], [S], M, N, P)` uses `NZMAX = 0`.
- `S = sparse([I], [J], [S], P)` uses `M = 0` and `N = 0`. This generates the ultimate sparse matrix, an M-by-N all zero matrix.
- `S = sparse(M, N, P)` abbreviates `sparse([I], [J], [S], M, N, 0, P)`. This also generates an M-by-N all zero matrix.

Remarks

The recommended method of creating a sparse dmat is with `spalloc`.

map/spalloc

Allocate space for a sparse dmat.

Syntax

```
S = spalloc(M, N, NZMAX, P)
```

Description

`S = spalloc(M, N, NZMAX, P)` creates an M-by-N all zero sparse dmat with room to eventually hold NZMAX nonzeros.

Remarks

NZMAX applies to the overall dmat, not individual processors. NZMAX is evenly distributed across processors

dmat/sparse

Converts a dmat to a sparse dmat

Syntax

```
S = sparse(X)
```

Description

`S = sparse(X)` converts a full dmat to sparse form by squeezing out any zero elements. If x is already a sparse dmat, `sparse(X)` returns X.

Data analysis and Fourier transforms

dmatrix/conv2

Two dimensional convolution.

Syntax

```
C = conv2(A, B, 'shape')
```

Description

`C = conv2(A, B, 'shape')` performs the 2D convolution of `dmatrix A` and `double B`. Returns a subsection of the 2D convolution with size specified by `'shape'`.

Remarks

Only `'shape' == 'same'` is supported, which returns the central part of the convolution of the same size as `A`.

dmatrix/fft

Discrete Fourier transform on a `dmatrix`.

Syntax


```
Y = fft(X)
Y = fft(X, [], DIM)
Y = fft(X, N, DIM)
```

Description

`fft(X)` is the discrete Fourier transform (DFT) of matrix `x`. The FFT operation is applied to each column. If the matrix `x` is row distributed, `fft` displays a warning and remaps `x`. Calls MATLAB `fft` on the local part.

`fft(X, [], DIM)` or `fft(X, N, DIM)` applies FFT across dimension `DIM`. `fft(X, N, DIM)` returns the `N`-point DFT. If `x` is distributed along a dimension other than dimension `DIM`, displays a warning and remaps `x` along the dimension `DIM`. Calls the MATLAB `fft` on the local part.

For example, suppose that `X` is distributed row-wise. Calling `Y = fft(X, [], 1)`, which will perform a FFT on each column, will perform the following remapping:



Remarks

`fft` supports 2D and 3D `dmats`.

If `fft` remaps `x`, `y` has the new map (different from the original map passed in with `x`).

Data types and structures

dmatrix/double

Converts each local part of the `dmatrix` to double precision floating point.

Syntax

```
D = double(X)
```

Description

`D = double(X)` returns the double precision value for the local part of `x` in a `dmatrix` `D` with the same mapping and dimensions as `x`. If the local part of `D` is already double precision, `double` has no effect.

dmatrix/single

Converts each local part of the `dmatrix` to single precision floating point.

Syntax

```
D = single(X)
```

Description

`D = single(X)` returns the single precision value for the local part of `x` in a `dmatrix` `D` with the same mapping and dimensions as `x`. If the local part of `D` is already single precision, `double` has no effect.

dmatrix/uint8, dmatrix/uint16, dmatrix/uint32, dmatrix/uint64

Converts each local part of the `dmatrix` to unsigned integer.

Syntax

```
D = uint8(X)
D = uint16(X)
D = uint32(X)
D = uint64(X)
```

Description

`D = uint*(X)` returns the unsigned integer value for the local part of `X` in a `dmatrix` `D` with the same mapping and dimensions as `X`. If the local part of `D` is already unsigned integer, `uint*` has no effect.

dmatrix/int8, dmatrix/int16, dmatrix/int32, dmatrix/int64

Converts each local part of the `dmatrix` to signed integer.

Syntax

```
D = int8(X)
D = int16(X)
D = int32(X)
D = int64(X)
```

Description

`D = int*(X)` returns the signed integer value for the local part of `X` in a `dmatrix` `D` with the same mapping and dimensions as `X`. If the local part of `D` is already signed integer, `int*` has no effect.

Signal Processing Toolbox

dmat/dct

Distributed discrete cosine transform

Syntax

$$B = \text{dct}(A)$$

Description

$B = \text{dct}(A)$ returns the discrete cosine transform of A . The vector B is the same size as A and contains the discrete cosine transform coefficients.

If A is a matrix, the `dct` operation is applied to each column. This transform can be inverted using `idct`.

dmat/idct

Distributed inverse discrete cosine transform

Syntax

$$A = \text{idct}(B)$$

Description

$A = \text{idct}(B)$ inverts the `dct` transform, returning the original vector if B was obtained using $B = \text{dct}(A)$.

If B is a matrix, the `idct` operation is applied to each column.

Index

dmat
 abs, 46
 acos, 39
 acosd, 40
 acosh, 41
 acot, 39
 acotd, 40
 acoth, 41
 acsc, 39
 acscd, 40
 acsch, 41
 agg, 28
 agg_all, 29
 angle, 46
 asec, 39
 asecd, 40
 asech, 41
 asin, 39
 asind, 40
 asinh, 41
 atan, 39
 atand, 40
 atanh, 41
 ceil, 48
 complex, 46
 conj, 47
 conv2, 58
 cos, 36
 cosd, 37
 cosh, 38
 cot, 36
 cotd, 37
 coth, 38
 csc, 36
 cscd, 37
 csch, 38
 dct, 62
 display, 17
 double, 60
 eq, 53
 exp, 42
 expm1, 42
 fft, 58
 find, 27
 fix, 48
 floor, 48
 global_block_range, 18
 global_block_ranges, 19
 global_ind, 20
 global_inds, 21
 global_range, 22
 global_ranges, 24
 gt, 54
 idct, 62
 imag, 47
 int16, 61
 int32, 61
 int64, 61
 int8, 61
 ldivide, 52
 local, 29
 log, 42

Index

log10, 43
log1p, 43
minus, 50
mtimes, 51
ndims, 16
plus, 50
pow2, 44
power, 52
put_local, 30
rdivide, 53
real, 47
reallog, 44
realpow, 44
realsqrt, 45
round, 49
sec, 36
secd, 37
sech, 38
sin, 36
sind, 37
single, 60
sinh, 38
size, 16
sparse, 57
sqrt, 45
subsasgn, 32
subsref, 33
synch, 31
tan, 36
tand, 37
tanh, 38
times, 51
transpose, 55
uint16, 61
uint32, 61
uint64, 61
uint8, 61
map
 display, 17
 eq, 54
 inmap, 26
 map, 9
 ne, 54
 ones, 14
 rand, 15
 spalloc, 57
 sparse, 56
 subsasgn, 34
 subsref, 34
 zeros, 13
MatMPI_Delete_all, 8
MPI_Abort, 7
MPI_Run, 8
pMATLAB, 6
pMatlab_Finalize, 7
pMatlab_Init, 6
pMatlab_ver, 7
remap, 31
transpose_grid, 35