

1 **A new adaptive data transfer library for model coupling**

2 **C. Zhang^{2,1}, L. Liu^{1,3}, G. Yang^{2,1,3}, R. Li^{2,1}, and B. Wang^{1,3,4}**

3 [1]{Ministry of Education Key Laboratory for Earth System Modeling, Center for Earth
4 System Science (CESS), Tsinghua University, Beijing, China}

5 [2]{Department of Computer Science and Technology, Tsinghua University, Beijing, China}

6 [3]{Joint Center for Global Change Studies (JCGCS), Beijing, China}

7 [4]{State Key Laboratory of Numerical Modelling for Atmospheric Sciences and Geophysical
8 Fluid Dynamics (LASG), Institute of Atmospheric Physics, Chinese Academy of Sciences,
9 Beijing, China.}

10 Correspondence to: L. Liu (liuli-cess@tsinghua.edu.cn), G. Yang (ygw@tsinghua.edu.cn)

11 **Abstract**

12 Data transfer means transferring data fields from a sender to a receiver. It is a fundamental
13 and most frequently used operation of a coupler. Most versions of state-of-the-art couplers
14 currently use an implementation based on the point-to-point (P2P) communication of the
15 Message Passing Interface (MPI) (refer such an implementation as “P2P implementation” for
16 short). In this paper, we revealed the drawbacks of the P2P implementation when the parallel
17 decompositions of the sender and the receiver are different, including low communication
18 bandwidth due to small message size, variable and big communication depth, as well as
19 network contention. To overcome these drawbacks, we proposed a butterfly implementation
20 for data transfer. Although the butterfly implementation can outperform the P2P
21 implementation in many cases, it degrades the performance when the sender and the receiver
22 have similar parallel decompositions or the number of processor cores used for running
23 models is small. To make data transfer always keep the optimal performance, we designed
24 and implemented an adaptive data transfer library that combines the advantages of both
25 butterfly implementation and P2P implementation. As the adaptive data transfer library can
26 adaptively use the better implementation for data transfer, it outperforms the P2P
27 implementation in many cases while does not decrease the performance in any cases. Now,
28 the adaptive data transfer library is open to the public and has been imported into a coupler
29

1 version C-Coupler1 for performance improvement of data transfer. We believe that other
2 couplers can also benefit from it.

3

4 **1 Introduction**

5 Climate System Models (CSMs) and Earth System Models (ESMs) are fundamental tools for
6 simulating, predicting and projecting climate. A CSM or an ESM generally integrates several
7 component models, such as an atmosphere model, a land surface model, an ocean model and a
8 sea-ice model, into a coupled system to simulate the behaviours of the climate system,
9 including the interactions between components of the climate system. More and more coupled
10 models have sprung up in the world. For example, the number of coupled model
11 configurations in the Coupled Model Intercomparison Project (CMIP) has increased from less
12 than 30 (used for CMIP3) to more than 50 (used for CMIP5).

13 High-performance computing is an essential technical support for model development,
14 especially for higher and higher resolutions of models. Modern high-performance computers
15 integrate an increasing number of processor cores for higher and higher computation
16 performance. Therefore, efficient parallelization, which enables a model to utilize more
17 processor cores for acceleration, becomes a technical focus in model development; and a
18 number of component models with efficient parallelization have sprung up. For example, the
19 Community Ice CodE (CICE; Hunke et al., 2008, 2013) at 0.1 °horizontal resolution can scale
20 to 30,000 processor cores on the IBM Blue Gene/L (Dennis et al., 2008); the Parallel Ocean
21 Program (POP; Kerbyson, 2005; Smith et al., 2010) at 0.1 °horizontal resolution can also
22 scale to 30,000 processor cores on the IBM Blue Gene/L and 10,000 processor cores on a
23 Cray XT3 (Dennis, 2007); the Community Atmosphere Model (CAM; Morrison et al., 2008;
24 Neale et al., 2010, 2012) with a spectral element dynamical core (CAM-SE) at 0.25 °
25 horizontal resolution can scale to 86,000 processor cores on a Cray XT5 (Dennis et al., 2012).

26 A coupler is an important component in a coupled system. It links component models together
27 to construct a coupled model, and controls the integration of the whole coupled model
28 (Valcke et al, 2012). A number of couplers now are available, e.g., the Model Coupling
29 Toolkit (MCT; Jacob et al., 2005), the Ocean Atmosphere Sea Ice Soil coupling software
30 (OASIS) coupler (Redler et al., 2010; Valcke, 2013; Valcke et al, 2015), the Earth System
31 Modelling Framework (ESMF; Hill et al., 2004), the CPL6 coupler (Craig et al., 2005), the
32 CPL7 coupler (Craig et al., 2012), the Flexible Modelling System (FMS) coupler (Balaji et al.,

1 2006), the Bespoke Framework Generator (BFG; Ford et al., 2006; Armstrong et al., 2009)
2 and the community coupler version 1 (C-Coupler1; Liu et al., 2014).

3 A coupler generally has much smaller overhead than the component models in current
4 coupled systems. However, it is potentially a time-consuming component in future coupled
5 models. This is because more and more component models (such as land-ice model,
6 chemistry model and biogeochemical model) will be coupled into a coupled model, and the
7 coupling frequency between component models will be higher and higher. Data transfer is a
8 fundamental and frequently used operation in a coupler. It is responsible for transferring data
9 fields between the processes of two component models and for rearranging data fields among
10 processes of the same component model for parallel data interpolation.

11 A coupler may become a bottleneck for efficient parallelization of future coupled models. The
12 most obvious reason is that the current implementation of data transfer in a state-of-the-art
13 coupler may be not efficient enough. For example, due to the low efficiency of data transfer,
14 the coupling from a component model with a horizontal grid (of 576×384 grid points) to
15 another component model with another horizontal grid (of 3600×2400 grid points) can only
16 scale to about 500 processor cores when using the CPL7 coupler (Craig et al., 2012).
17 Therefore, it is highly desirable to improve the parallel data transfer of couplers.

18 In this study, we first propose a butterfly implementation of data transfer. Since the P2P
19 implementation and the butterfly implementation can outperform each other in different cases
20 (Sect. 5), we next develop an adaptive data transfer library that includes both implementations
21 and can adaptively use the better one for data transfer. Performance evaluation demonstrates
22 that such a library significantly outperforms the P2P implementations in most cases and does
23 not degrade the performance in any case. This library has been imported into C-Coupler1 with
24 slight code modification. We believe that other couplers can also benefit from it.

25 The remainder of this paper is organized as follows. We briefly introduce the implementation
26 of data transfer in existing couplers in Section 2. Details of the butterfly implementation and
27 the adaptive data transfer library are presented in Sections 3 and 4, respectively. The
28 performances of data transfer implementations are evaluated in Section 5. Conclusions are
29 given in Section 6.

1 **2 Data transfer implementations in existing couplers**

2 **2.1 P2P implementation**

3 Almost all state-of-the-art couplers use a similar implementation for data transfer. To achieve
4 parallel data transfer, MCT first generates a communication router (known as the data
5 mapping between processes) according to the parallel decompositions (the distribution of grid
6 points among the processes) of the sender and the receiver, and then uses the point-to-point
7 (P2P) communication of the Message Passing Interface (MPI) to transfer the data. A data
8 field will be transferred from a process of the sender to a process of the receiver, only when
9 the two processes have common grid points. In the following context, we call this “P2P
10 implementation” for short.

11 Since MCT has already been imported into OASIS3-MCT, the CPL6 coupler and the CPL7
12 coupler, these couplers also use the P2P implementation for data transfer. Although the other
13 couplers such as ESMF, OASIS4, the FMS coupler and C-Coupler1 do not directly import
14 MCT, they also use the P2P implementation for data transfer.

15 **2.2 Performance bottlenecks of the P2P implementation**

16 To motivate this work, we first investigate the performance characteristics of the P2P
17 implementation, and therefore derive a benchmark from a real coupled model GAMIL2-
18 CLM3, which includes GAMIL2 (Li et al., 2013) that is an atmosphere model and CLM3
19 (Oleson et al., 2004; Dickinson et al., 2006) that is a land surface model. GAMIL2 and CLM3
20 share the same horizontal grid of 7,680 (128×60) grid points, but have different parallel
21 decompositions: GAMIL2 uses a regular 2-D parallel decomposition, while CLM3 uses an
22 irregular 2-D parallel decomposition where the grid points are assigned to the processes in a
23 round-robin fashion.

24 In this benchmark, there is only the data transfer with the P2P implementation between the
25 sender and the receiver with the same horizontal grid of GAMIL2-CLM3. The parallel
26 decomposition of the sender is derived from CLM3, and the parallel decomposition of the
27 receiver is derived from GAMIL2. A high-performance computer named Tansuo100 at
28 Tsinghua University, China is used for the performance tests. It has 700 computing nodes,
29 each of which contains two six-core Intel Xeon X5670 CPUs and 32 GB main memory. All

1 computing nodes are connected by a high-speed InfiniBand network with peak
2 communication bandwidth of 5 GB/s.

3 To evaluate the parallel performance of the P2P implementation, 14 2-D coupling fields are
4 transferred between the sender and the receiver. In each test, the sender and the receiver use
5 the same number of processes. Since there are 12 processor cores on each computing node,
6 the number of processes is set to be an integral multiple of 12. When the number of processes
7 is less than 12, the sender and the receiver are located on two different computing nodes. The
8 sender and the receiver do not share the same computing node, so the communication of the
9 P2P implementation must go through the InfiniBand network.

10 Figure 1 demonstrates that poor parallel scalability of the P2P implementation can be
11 obtained when the parallel decompositions of the sender and receiver are different. It is well
12 known that the communication performance heavily depends on message size. As shown in
13 Fig. 2, the P2P communication bandwidth achieved generally increases with message size. So
14 when the message size is small (for example, smaller than 4 KB), the communication
15 bandwidth achieved is very low. The message size in the P2P implementation decreases with
16 the increment of number of processes of models (Fig. 3), indicating that the communication
17 bandwidth becomes lower with the increment of number of processes. The performance of
18 data transfer also heavily depends on another term of communication depth, which is defined
19 as the number of the communications that a process is associated with. The communication
20 depth is determined by the parallel decompositions of the sender and the receiver. In the P2P
21 implementation, if one process of the sender/receiver has common grid points with N
22 processes of the receiver/sender, the communication depth of this process is N . As shown in
23 Fig. 4, the variation of average communication depth in the P2P implementation is consistent
24 with the variation of the execution time of the P2P implementation in Fig. 1: both the average
25 communication depth and the execution time of the P2P implementation increase with the
26 number of cores from 6 to 48, and go down with the number of cores from 96 to 192. Lower
27 execution time of the P2P implementation will be obtained if more cores are used (the
28 maximum number of cores in both Fig. 1 and Fig. 4 is limited to 192 because GAMIL2-
29 CLM3 will not be further accelerated when using more cores) since the average
30 communication depth will further go down. To further reveal possible reasons for the poor
31 parallel scalability, we evaluate the ideal performance and actual performance in Fig. 5. The
32 ideal performance is much better than the actual performance, and the ratio between the ideal

1 performance and the actual performance significantly increases with the increment of the
2 number of processes. The significant gap between the ideal performance and the actual
3 performance is due to network contention. For example, when multiple P2P communications
4 share the same sender process or receiver process, they must wait in order.

5 **3 Butterfly implementation for better performance of data transfer**

6 The drawbacks of the P2P implementation when the sender and the receiver use different
7 parallel decompositions can be identified as low communication bandwidth due to small
8 message size, variable and big communication depth, as well as network contention. To
9 overcome these drawbacks, a prospective solution is to organize the transfer of data using a
10 better algorithm, e.g., the butterfly algorithm (Fig. 6), which has already been studied in
11 computing sciences (Chong et al., 1994; Foster, 1995; Heckbert et al., 1995; Hemmert et al.,
12 2005; Kim et al., 2007; Jan et al., 2013; Petagon et al, 2016). In hardware aspect, the
13 traditional butterfly algorithm and its transformation have been used to design networks
14 (Chong et al., 1994; Kim et al., 2007); in software aspect, the butterfly algorithm has been
15 used to improve the parallel algorithms with all-to-all communications (Foster, 1995), e.g.,
16 Fast Fourier Transform (FFT; Heckbert et al., 1995; Hemmert et al., 2005), matrix
17 transposition (Petagon et al, 2016) and sorting (Jan et al., 2013).

18 Unfortunately, the classical butterfly algorithm cannot be used as is to improve data transfer,
19 because it requires that one process communicates with every other process, that the
20 communication load among processes is balanced and that the number of processes must be a
21 power of 2. In practice, data transfer for model coupling has different characteristics, i.e., one
22 process needs to communicate with a part of other processes, the communication load among
23 processes is always unbalanced and the number of processes cannot be restricted to a power
24 of 2. Therefore, we propose here a new implementation of data transfer involving an
25 additional butterfly kernel to transfer data from the sender with the source parallel
26 decomposition to the receiver with the target parallel decomposition. As the number of
27 processes of the butterfly kernel must be a power of 2, while the number of processes of the
28 sender or the receiver are not necessarily, the butterfly kernel has its own source parallel
29 decomposition and target parallel decomposition, and process mappings are needed from the
30 sender onto the butterfly kernel and from the butterfly kernel onto the receiver (see Fig. 7).
31 Next, we present the butterfly kernel and the process mappings, respectively.

1 3.1 Butterfly kernel

2 The first question for the butterfly kernel is how to decide its number of processes. Any
3 process of the sender or receiver can be used as a process for the butterfly kernel. Given that
4 the total number of unique processes of the sender and receiver is N_T , the number of processes
5 of the butterfly kernel (N_B) can be any power of 2, which is no larger than N_T . We propose to
6 select the maximum number in order for maximum utilization of resources. We prefer to pick
7 out unique processes first from the sender, and then from the receiver if the sender does not
8 have enough processes.

9 The butterfly kernel is responsible for rearranging the distribution of data among the
10 processes from the source parallel decomposition to the target parallel decomposition. Given
11 the number of processes $N=2^n$, there are n stages in the butterfly kernel. In a stage, all
12 processes are divided into a number of pairs and the two processes of a pair uses MPI P2P
13 communication to exchange data. After each stage, the number of butterfly kernel processes
14 that may have the data that will finally belong to any one process on the target parallel
15 decomposition will become a half. Figure 6 is an example for further illustration, where D^i_j
16 means the data is originally in process P_i according to the source parallel decomposition and
17 is finally in process P_j according to the target parallel decomposition. Before the first stage,
18 all processes ($P_0\sim P_7$) may have the data of P_0 on the target parallel decomposition. After the
19 first stage, only four processes (P_0, P_2, P_4 and P_6) may have that data; and after the second
20 stage, only two processes (P_0 and P_4) may have it.

21 To reveal the advantages and disadvantages of the two implementations, we measure the
22 characteristics of the two implementations based on the benchmark introduced in Section 2.2.
23 The results show that the total message size transferred by the butterfly implantation is larger
24 than that by the P2P implementation (Fig. 8), which is the major disadvantage of the butterfly
25 implementation. Meanwhile, comparing with the P2P implementation, the butterfly
26 implementation can have the following advantages:

- 27 1) bigger message size for better communication bandwidth (Fig. 9);
- 28 2) balanced and smaller communication depth among processes (Fig. 10);
- 29 3) ordered communications among processes and fewer communications operated
30 concurrently (Fig. 10), which can dramatically reduce network contention.

1 3.2 Process mapping

2 In this subsection, we will introduce the process mappings from the sender to the butterfly
3 kernel and from the butterfly kernel to the receiver. To minimize the overhead of process
4 mapping from the butterfly kernel to the receiver, we map one or multiple processes of the
5 butterfly kernel onto a process of the receiver if the butterfly kernel has more processes than
6 the receiver; otherwise, we map a process of the butterfly kernel onto one or multiple
7 processes of the receiver. In other words, there is no multiple-to-multiple process mapping
8 between the butterfly kernel and the receiver. Similarly, there is no multiple-to-multiple
9 process mapping between the sender and the butterfly kernel.

10 Processes of the sender or the receiver may be unbalanced in terms of the size of the data
11 transferred, which may result in unbalanced communications among processes of the butterfly
12 kernel. As mentioned in Section 3.1, at each stage of the butterfly kernel, all processes are
13 divided into a number of pairs, each of which is involved in P2P communications. To
14 improve the balance of communications among the processes in the butterfly kernel, one
15 solution is to try to make the process pairs at each stage more balanced in terms of data size of
16 P2P communications, so we propose to reorder the processes of the sender or the receiver
17 according to data size. At the first stage, each time we pick out the process with the largest
18 data size and the process with the smallest data size from the remaining processes that have
19 not been paired, to generate a process group. For the next stage, the outputs of two process
20 groups from the previous stage are paired into a bigger process groups in a similar way. After
21 finishing the iterative pairing throughout all stages, all processes of the sender or the receiver
22 are reordered.

23 The iterative pairing also requires the number of processes to be a power of 2. Given that the
24 number of processes of the sender (or receiver) is N_C and the number of processes of the
25 butterfly kernel is N_B , we first pad empty processes (whose data size is zero) before the
26 iterative pairing to make the number of processes of the sender (or receiver) be a power of 2
27 (denoted N_P), which is no smaller than N_B . Therefore, the reordered N_P processes after the
28 iterative pairing can be divided into N_B groups, each of which contains N_P/N_B processes with
29 consecutive reordered indexes and maps onto a unique process of the butterfly kernel.

30 Figure 11 shows an example of the process mapping, where the sender has five processes (S_0 -
31 S_4 in Fig. 11a), the receiver has 10 processes (R_0 - R_9 in Fig. 11b), and the butterfly kernel uses
32 eight processes (B_0 - B_7 in Fig. 11c). At first, empty processes are padded to the sender (S_5 - S_7

1 in Fig. 11a) and the receiver (R_{10} - R_{15} in Fig. 11b). Next, the iterative pairing is conducted for
2 the sender and the receiver, respectively. The iterative pairing has three stages for the sender.
3 At the first stage, the eight processes of the sender are divided into four groups $\{S_1, S_7\}$,
4 $\{S_0, S_6\}$, $\{S_2, S_5\}$ and $\{S_4, S_3\}$ (Fig. 11a), according to the data size corresponding to each
5 process. These four process groups are divided into two bigger groups ($\{\{S_4, S_3\}, \{S_2, S_5\}\}$ and
6 $\{\{S_1, S_7\}, \{S_0, S_6\}\}$ at the second stage (Fig. 11a). Finally, one process group
7 $\{\{\{S_4, S_3\}, \{S_2, S_5\}\}, \{\{S_1, S_7\}, \{S_0, S_6\}\}\}$ is obtained at the third stage (Fig. 11a), and the eight
8 processes of the sender are reordered as $S_4, S_3, S_2, S_5, S_1, S_7, S_0$ and S_6 , each one being mapped
9 onto one process of the butterfly kernel (Fig. 11c). Similarly, the iterative pairing has four
10 stages for the receiver, and the 16 processes of the receiver are reordered as $R_9, R_{15}, R_7, R_{12},$
11 $R_4, R_8, R_3, R_{10}, R_1, R_{14}, R_5, R_{13}, R_0, R_6, R_2$ and R_{11} finally, with pairs of these being mapped
12 onto one process of the butterfly kernel (Fig. 11c).

13 **4 Adaptive data transfer library**

14 Now, we have two kinds of implementations (the P2P implementation and the butterfly
15 implementation) for data transfer. Although the butterfly implementation can effectively
16 improve the performance of data transfer in many cases (examples are given in Section 5), it
17 has some drawbacks: 1) it generally has a larger total message size than the P2P
18 implementation; 2) its number of stages is $\log_2 N$ (where N is the number of processes for the
19 butterfly kernel) (Foster, 1995), which may be bigger than the average communication depth
20 in the P2P implementation in some cases (for example, when the sender and the receiver use
21 the similar parallel decompositions). Therefore, it is possible that the P2P implementation
22 outperforms the butterfly implementation in some cases. To achieve optimal performance for
23 data transfer, we propose an adaptive data transfer library that can take the advantages of the
24 two implementations in all cases.

25 As introduced in Section 3.1, the butterfly implementation is divided into multiple stages.
26 Actually, the data transfer in one stage can be viewed as a P2P implementation with only one
27 MPI message per process. Inspired by this fact, we try to design an adaptive approach that can
28 combine the butterfly and P2P implementations, where some stages in the butterfly
29 implementation are skipped and replaced by P2P communication of more MPI messages per
30 process. When all stages of the butterfly implementation are skipped, the adaptive data
31 transfer library completely switches to the original P2P implementation. That is to say, the
32 adaptive data transfer can adaptively choose the optimal implementation from the P2P

1 implementation and the butterfly implementation. Figure 12 shows an example of the
2 adaptive data transfer library with eight processes, where Stage 2 of the butterfly
3 implementation is skipped and replaced by P2P communication of three MPI messages per
4 process.

5 The most significant challenge of such an adaptive approach is to determine which stage(s) of
6 the butterfly implementation should be skipped. The first attempt was to design a cost model
7 that can accurately predict the performance of data transfer in various implementations. We
8 eventually gave up this approach as it was almost impossible to accurately predict the
9 performance of the communications on a high-performance computer, especially when a lot
10 of users share the computer to run various applications. Performance profiling which means
11 directly measuring the performance of data transfer is more practical to determine an
12 appropriate implementation, because the simulation of earth system modelling always takes a
13 long time to run. Figure 13 shows our flowchart of how the adaptive data transfer library
14 determines an appropriate implementation. It consists of an initialization segment and a
15 profiling segment. The initialization segment generates the process mappings and a candidate
16 implementation that is a butterfly implementation with no skipped stages. The profiling
17 segment iterates through each stage of the butterfly implementation to determine whether the
18 current stage should be skipped or kept. In an iteration, the profiling segment first generates a
19 temporary implementation based on the candidate implementation where the current stage is
20 skipped, and then runs the temporary implementation to get the time the data transfer takes.
21 When the temporary implementation is more efficient than the candidate implementation, the
22 current stage is skipped and the temporary implementation replaces the candidate
23 implementation. When the profiling segment finishes, the appropriate implementation is set to
24 be the candidate implementation. To reduce the overhead introduced by the adaptive data
25 transfer library, the profiling segment truly transfers the data for model coupling. In other
26 words, before obtaining an optimal implementation, the data is transferred by the profiling
27 segment.

28 **5 Performance evaluation**

29 In this section, we empirically evaluate the adaptive data transfer library, through comparing
30 it to the P2P implementation and the butterfly implementation. Both toy models and realistic
31 models (GAMIL2-CLM3 and CESM) are used for the performance evaluation. GAMIL2-
32 CLM3 has been introduced in Section 2.2. CESM (Hurrell et al., 2013) is a state-of-the-art

1 ESM developed by the National Center for Atmospheric Research (NCAR). All the
2 experiments are run on the high performance computer Tansuo100.

3 Next, we will evaluate the overhead of initialization, the performance of transferring data
4 fields between two different component models and the performance of rearranging data
5 fields intra a component model for parallel interpolation.

6 **5.1 Overhead of initialization**

7 We first evaluate the initialization overhead of data transfer implementations. As shown in
8 Fig. 14, the initialization overhead of each implementation increases with the increment of
9 core number. The initialization overhead of the butterfly implementation is a little higher than
10 that of the P2P implementation, while the initialization overhead of the adaptive data transfer
11 library is 2-3 folds higher than that of the P2P implementation, because the adaptive data
12 transfer library uses extra time on the performance profiling (see Section 4). Considering that
13 one data transfer instance should only be initialized at the beginning and executed many times
14 in a coupled model, we can conclude that the initialization overhead of the adaptive data
15 transfer library is reasonable, especially when the simulation is executed for a very long time.

16 **5.2 Performance of data transfer between toy models**

17 The factors that can impact the performance of a data transfer implementation generally
18 include the communication depth, the size of the data to be transferred (also referred to as the
19 number of fields in this evaluation) and the number of cores used. In this subsection, we
20 evaluate the impact of each factor on the performance of data transfer for different
21 implementations. We first build two toy models that both use the same logically rectangular
22 grid of 192×480 grid points. Coupling fields are transferred between the two toy models. For
23 any test, the two toy models use the same number of cores. Next, we evaluate the
24 performance of data transfer through varying one factor while fixing the other two factors.

25 In the first experiment, we fix the number of cores to be 1024 and the number of coupling
26 fields to be 10, while only vary the communication depth in the P2P implementation. In each
27 test, all processes of the sender have the same communication depth. As the communication
28 depth is determined by the parallel decompositions of the sender and the receiver, we design
29 an algorithm (Algorithm 1) that can generate the parallel decompositions of the two toy
30 models according to the average communication depth of the sender in the P2P

1 implementation. Figure 15 shows the execution time of one data transfer with different
2 implementations when increasing the communication depth per sender process in the P2P
3 implementation from 1 to 90. The P2P implementation can outperform the butterfly
4 implementation when the communication depth is small (say, smaller than 12 in Fig. 15),
5 while the butterfly implementation can outperform the P2P implementation when the
6 communication depth is big (say, bigger than 12 in Fig. 15). The adaptive data transfer library
7 can adaptively choose the optimal implementation from the P2P implementation and the
8 butterfly implementation, and moreover, it improves the performance based on the butterfly
9 implementation when the communication depth is big, because some butterfly stages of the
10 butterfly implementation are skipped. When the communication depth is 90, the adaptive data
11 transfer library can achieve a 19.2-fold performance speedup compared to the P2P
12 implementation.

13 In the second experiment, we fix the number of cores and the communication depth per
14 sender process in the P2P implementation, and vary the number of coupling fields transferred.
15 Figure 16 shows the execution time of one data transfer with different implementations in this
16 experiment. The results show that the execution time of each implementation increases with
17 the increment of data size. When the communication depth per sender process in the P2P
18 implementation is small (Figs. 16a and 16b), the performance of the butterfly implementation
19 is poorer than that of the P2P implementation, especially when the number of 2-D coupling
20 fields gets bigger. When the communication depth per sender process in the P2P
21 implementation is big (Figs. 16c and 16d), the butterfly implementation significantly
22 outperforms the P2P implementation, however, the advantage of the butterfly implementation
23 decreases with the increment of the number of coupling fields. The results also demonstrate
24 that the adaptive data transfer library can adaptively choose the optimal implementation from
25 the P2P implementation and the butterfly implementation, and can further improve the
26 performance based on the butterfly implementation.

27 In the third experiment, we fix the communication depth per sender process in the P2P
28 implementation to be 24 and the number of coupling fields transferred to be 10, and vary the
29 number of cores used. Figure 17 shows the execution time of one data transfer with different
30 implementations when varying the number of cores. The P2P implementation outperforms the
31 butterfly implementation, when small number of cores are used (say, smaller than 256 in Fig.
32 17); while the butterfly implementation outperforms the P2P implementation when large

1 number of cores are used (say, larger than 256 in Fig.17). Similar to above two experiments,
2 the adaptive data transfer library can adaptively choose the optimal implementation from the
3 P2P implementation and the butterfly implementation.

4 The resolutions of models become higher and higher these days. How about the performance
5 of the data transfer implementations when model resolutions become higher? Higher model
6 resolutions mean that a model will use more processor cores for accelerating a simulation,
7 while the average number of grid points per processor core can remain constant. Considering
8 that the numbers of grid points are always balanced among the processes of a model, we make
9 each process (which runs on a unique processor core) of the toy models evenly have around
10 96 grid points in this evaluation, while enabling processes to have different communication
11 depth and different message sizes (the average communication depth of the sender in P2P
12 implementation is 34). As shown in Fig. 18, although the execution times of all data transfer
13 implementations increase when increasing the number of processor cores (from 64 to 1024),
14 the butterfly implementation significantly outperforms the P2P implementation. So the
15 adaptive data transfer library adaptively chooses the butterfly implementation, and further
16 slightly outperforms the butterfly implementation when each model uses more than 512 cores
17 because some butterfly stages are skipped.

18 **5.3 Performance of data transfer between realistic models**

19 In this subsection, we evaluate the performance using two realistic models: GAMIL2-CLM3
20 (horizontal resolution of $2.8^\circ \times 2.8^\circ$) and CESM (resolution of $1.9 \times 2.5_{\text{gx1v6}}$).

21 For CESM, we use the data transfer between the coupler CPL7 (Craig et al., 2012) and the
22 land surface model CLM4 (Oleson et al., 2004), where 32 2-D coupling fields on the CLM4
23 horizontal grid (the grid size is $144 \times 96 = 13824$) are transferred. Figure 19 shows the
24 performance of one data transfer of different implementations when increasing the number of
25 processes of both CPL7 and CLM4 from 6 to 192. When the number of processes is small
26 (say, smaller than 24 in Fig. 19), the butterfly implementation is much poorer than the P2P
27 implementation. In this case, the adaptive data transfer library chooses the P2P
28 implementation as the optimal implementation. However, when the number of processes gets
29 bigger (say, larger than 24 in Fig. 19), the butterfly implementation outperforms the P2P
30 implementation. In this case, the adaptive data transfer library based on the butterfly
31 implementation skips some stages, so it outperforms the butterfly implementation. Figure 19

1 also shows that the butterfly implementaion and the adaptive transfer library seem to
2 converge when increasing the number of cores per model. When each model uses 192 cores,
3 the adaptive data transfer library is 4.01 times faster than the P2P implementation.

4 For GAMIL2-CLM3, we use the data transfer from CLM3 to GAMIL2 where 14 2-D
5 coupling fields on the GAMIL2 horizontal grid (whose grid size is $128 \times 60 = 7680$) are
6 transferred. Figure 20 shows the execution time of one data transfer of each implementation
7 when increasing the number of processes of both GAMIL2 and CLM3 from 6 to 192. The
8 results in Fig. 20 confirm that the adaptive data transfer library can adaptively choose the
9 optimal implementation from the P2P implmentation and the butterfly implementation.
10 Compared to the P2P implementation, the adaptive data transfer library achieves an 11.68-
11 fold performance speedup when the number of processes is 96, but achieves a much lower
12 speedup (only 3.48-fold) when the number of processes is 192. This is because the average
13 communication depth per process in the P2P implementation reduces from 32 to 18 when the
14 number of process increases from 96 to 192.

15 **5.4 Performance of data rearrangement for interpolation**

16 Besides data transfer between different component models, there is another kind of data
17 transfer in model coupling that rearranges data inside a model for parallel interpolation of
18 fields between different grids. Here, we use the data rearrangement for the parallel
19 interpolation from the atmosphere grid (whose grid size is $144 \times 96 = 13824$) to the ocean grid
20 (whose grid size is $320 \times 384 = 122880$) in the coupled model CESM for further evaluation. As
21 shown on Fig. 21, the P2P implementation significantly outperforms the butterfly
22 implementation. This is because the corresponding parallel decompositions for data
23 rearrangement are always similar while similar parallel decompositions generally introduce
24 small communication depth. For example, average communication depth in the P2P
25 implementation corresponding to Fig. 21 is only 6.49 when the model uses 96 cores. In this
26 case, the P2P implementation is chosen as the optimal implementation of the data transfer
27 library, so the data transfer library library does not provide real benefit compared to the P2P
28 implementation.

1 **5.5 Performance improvement for a coupled model**

2 With the performance improvement of data transfer, we expect that the adaptive data transfer
3 library will improve the performance of coupled models. For this evaluation, we first
4 imported the adaptive data transfer library into C-Coupler1 and then used the coupled model
5 GAMIL2-CLM3 that uses C-Coupler1 for coupling to measure performance results. As
6 shown in Fig. 22, the adaptive data transfer library achieves higher performance improvement
7 (when the P2P implementation is used as the baseline) for GAMIL2-CLM3 when using more
8 processor cores. When each component model uses 128 processor cores, the butterfly
9 implementation achieves ~4.6% performance improvement, and the adaptive data transfer
10 library achieves ~6.9% performance improvement. So the data transfer library can improve
11 the performance of data transfer, and then improve the performance of the whole coupled
12 models.

13 **6 Conclusions**

14 Data transfer is a fundamental and most frequently used operation in a coupler. This paper
15 showed that the P2P implementation currently used in most state-of-the-art couplers for data
16 transfer is inefficient when the parallel decompositions of the sender and the receiver are
17 different, and further revealed the corresponding performance bottlenecks. To overcome these
18 bottlenecks, we proposed the butterfly implementation that can outperform the P2P
19 implementation in many cases, however, degrades the performance in some cases, for example,
20 when a small number of cores are used to run models or the parallel decompositions of the
21 sender and receiver are similar. We therefore further designed and implemented an adaptive
22 data transfer library that can not only adaptively choose an optimal implementation from the
23 P2P implementation and the butterfly implementation, but also further improve the performance
24 based on the butterfly implementation through skipping some butterfly stages. Compared to
25 the P2P implementation, the adaptive data transfer library can improve the performance of
26 data transfer when the parallel decompositions of the sender and the receiver are different.

27 The initialization overhead for the adaptive data transfer library could become expensive
28 when using a large number of processor cores. In the future version, the adaptive data transfer
29 will allow users to record the results of performance profiling offline to save the time used for
30 performance profiling in next runs of the same coupled model.

31 **Code availability**

1 The source code of the adaptive data transfer library version 1.0 is available at
2 https://github.com/zhang-cheng09/Data_transfer_lib.

3 **Acknowledgements**

4 This work is supported in part by the Natural Science Foundation of China (no. 41275098),
5 the National Grand Fundamental Research 973 Program of China (no. 2013CB956603) and
6 the Tsinghua University Initiative Scientific Research Program (no. 20131089356).

7

1 **References**

- 2 Armstrong, C. W., Ford, R. W., and Riley, G. D.: Coupling integrated Earth System Model
3 components with BFG2, *Concurrency and Computation: Practice and Experience*,
4 2009;21;767–791, doi:10.1002/cpe.1348, 2009.
- 5 Balaji, V., Anderson, J., Held, I., Winton, M., Durachta, J., Malyshev, S., and Stouffer, R. J.:
6 The Exchange Grid: a mechanism for data exchange between Earth system components on
7 independent grids, In *Parallel Computational Fluid Dynamics 2005 Theory and Applications*,
8 2006, 179-186, doi: 10.1016/B978-044452206-1/50021-5, 2006.
- 9 Chong, F. T., and Brewer, E. A.: Packaging and multiplexing of hierarchical scalable
10 expanders, *Parallel Computer Routing and Communication*, Springer Berlin Heidelberg,
11 1994:200-214.
- 12 Craig, A. P., Vertenstein, M., and Jacob, R.: A new flexible coupler for Earth system
13 modelling developed for CCSM4 and CESM1, *Int. J. High Perform. C.*, 26, 31-42,
14 doi:10.1177/1094342011428141, 2012.
- 15 Craig, A. P., Jacob, R., Kauffman, B., Bettge, T., Larson, J., Ong, E., Ding, C., and He, Y.:
16 CPL6: the New Extensible, High Performance Parallel Coupler for the Community Climate
17 System Model, *Int. J. High Perform. C.*, 19, 309–327, 2005.
- 18 Dennis, J. M.: Inverse space-filling curve partitioning of a global ocean model, In *IEEE*
19 *International Parallel & Distributed Processing Symposium*, Long Beach, CA, 2007.
- 20 Dennis, J. M. and Tufo, H. M.: Scaling climate simulation applications on the IBM Blue
21 Gene/L system, *IBM J. Res. Dev.*, 52, 117-126, DOI:10.1147/rd.521.0117, 2008.
- 22 Dennis, J. M., Edwards, J., Evans, K. J., Guba, O., Lauritzen, P. H., Mirin, A. A., St-Cyr, A.,
23 Taylor, M. A., and Worley, P. H.: CAM-SE: a scalable spectral element dynamical core for
24 the Community Atmosphere Model, *Int. J. High Perform. C.*, 26, 74-89,
25 doi:10.1177/1094342011428142, 2012.
- 26 Dickinson, R. E., Oleson, K. W., Bonan, G., Hoffman, F., Thornton, P., Vertenstein, M.,
27 Yang, Z.-L., and Zeng X.: The Community Land surface model and its climate statistics as a
28 component of the Community Climate System Model, *Journal of Climate*, 19(11), 2302–2324,
29 2006.

1 Ford, R. W., Riley, G. D., Bane, M. K., Armstrong, C. W., and Freeman, T. L.: GCF: a
2 general coupling framework, *Concurrency and Computation: Practice and Experience*, 18(2),
3 163–181, 2006.

4 Foster I.: *Designing and building parallel programs: concepts and tools for parallel software*
5 *engineering*, Addison-Wesley, 1995.

6 Heckbert P.: *Fourier Transforms and the Fast Fourier Transform (FFT) Algorithm*, *Computer*
7 *Graphics*, 2: 15-463, 1995.

8 Hemmert, K. S., and K. D. Underwood.: *An analysis of the double-precision floating-point*
9 *FFT on FPGAs*. *Field-Programmable Custom Computing Machines*, 2005. FCCM 2005. 13th
10 *Annual IEEE Symposium on IEEE*, 2005:171-180.

11 Hill, C., DeLuca, C., Balaji, V., Suarez, M., and da Silva, A.: *The Architecture of the Earth*
12 *System Modelling Framework*, *Computing in Science & Engineering*, 6(1), 18–28, 2004.

13 Hurrell, J. W., Holland, M. M., Gent, P. R., Ghan, S., Kay, J. E., Kushner, P. J., Lamarque, J.-
14 F., Large, W. G., Lawrence, D., Lindsay, K., Lipscomb, W. H., Long, M. C., Mahowald, N.,
15 Marsh, D. R., Neale, R. B., Rasch, P., Vavrus, S., Vertenstein, M., Bader, D., Collins, W. D.,
16 Hack, J. J., Kiehl, J., and Marshall, S.: *The Community Earth System Model: a framework for*
17 *collaborative research*, *Bulletin of the American Meteorological Society*, 94(9), 1339–1360,
18 2013.

19 Hunke, E. C. and Lipscomb W. H.: *CICE: the Los Alamos Sea Ice Model Documentation and*
20 *Software User’s Manual 4.0*, Technical Report LA-CC-06-012, Los Alamos National
21 Laboratory, T-3 Fluid Dynamics Group, 2008.

22 Hunke, E. C., Lipscomb, W. H., Turner, A. K., Jeffery, N., and Elliott, S.: *CICE: the Los*
23 *Alamos Sea Ice Model Documentation and Software User’s Manual Version 5.0*, LA-CC-06-
24 012, Los Alamos National Laboratory, Los Alamos NM, 87545, 115, 2013.

25 Jacob, R., Larson, J., and Ong, E.: *M × N Communication and Parallel Interpolation in*
26 *Community Climate System Model version 3 using the Model Coupling Toolkit*, *International*
27 *Journal of High Performance Computing Applications*, 19(3), 293–307, 2005.

28 Jan, B., Montrucchio, B., Ragusa, C., Khan, F. G., and Khan, O.: *Parallel butterfly sorting*
29 *algorithm on gpu*, Acta Press, 2013.

1 Kerbyson, D. J., and Jones, P. W.: A performance model of the parallel ocean program,
2 International Journal of High Performance Computing Applications, 19(3), 261-276,
3 doi:10.1177/1094342005056114, 2005.

4 Kim J., Dally W. J., and Abts D.: Flattened butterfly: A cost-efficient topology for high-radix
5 networks, ISCA, 2007, 35(2):126-137.

6 Li, L. J., Wang, B., Dong, L., Liu, L., Shen, S., Hu, N., Sun, W., Wang, Y., Huang, W., Shi,
7 X., Pu, Y., G. and Yang.: Evaluation of Grid-point Atmospheric Model of IAP LASG version
8 2 (GAMIL2), Advances in Atmospheric Sciences, 30, 855–867, doi:10.1007/s00376-013-
9 2157-5, 2013.

10 Liu, L., Yang, G., Wang, B., Zhang, C., Li, R., Zhang, Z., Ji, Y., and Wang, L.: C-Coupler1: a
11 Chinese community coupler for Earth system modeling, Geoscientific Model Development,
12 7(5), 2281-2302, doi:10.5194/gmd-7-2281-2014, 2014.

13 Morrison, H., and A. Gettelman: A new two-moment bulk stratiform cloud microphysics
14 scheme in the Community Atmosphere Model, version 3 (CAM3). Part I: Description and
15 numerical tests, Journal of Climate, 21(15), 3642–3659, doi:10.1175/2008JCLI2105.1, 2008.

16 Neale, R. B., Richter, J. H., Conley, A. J., Park, S., Lauritzen, P. H., Gettelman, A.,
17 Williamson, D. L., Rasch, P. J., Vavrus, S. J., Taylor, M. A., Collins, W. D., Zhang, M., and
18 Lin, S.: Description of the NCAR Community Atmosphere Model (CAM 4.0), National
19 Center for Atmospheric Research Ncar Koha Openpat, TN-485+STR, 222p., 2010.

20 Neale, R. B., Chen, C. C., Gettelman, A., Lauritzen, P. H., Park, S., Williamson, D. L.,
21 Conley, A. J., Garcia, R., Kinnison, D., Lamarque, J. F., Marsh, D., Mills, M., Smith, A. K.,
22 Tilmes, S., Vitt, F., Morrison, H., Cameron-Smith, P., Collins, W. D., Iacono, M. J., Easter, R.
23 C., Ghan, S. J., Liu, X., Rasch, P. J., and Taylor, M. A.: Description of the NCAR
24 Community Atmosphere Model (CAM 5.0), National Center for Atmospheric Research Ncar
25 Koha Openpat, TN-486+STR, 289p., 2012

26 Oleson, K. W., Dai, Y., Bonan, G., Bosilovich, M., Dickinson, R., Dirmeyer, P., Hoffman, F.,
27 Houser, P., Levis, S., Niu, G. Y., Thornton, P., Vertenstein, M., Yang, Z. L., and Zeng, X.:
28 Technical Description of the Community Land Surface Model (CLM), National Center for
29 Atmospheric Research Ncar Koha Openpat, TN-461+STR, 186p., 2004.

1 Petagon, R., and Werapun, J.: Embedding the optimal all-to-all personalized exchange on
2 multistage interconnection networks + + mathContainer Loading Mathjax, *Journal of Parallel
3 & Distributed Computing* 88(2016):16-30.

4 Redler, R., Valcke, S., and Ritzdorf, H.: OASIS4—a coupling software for next generation
5 Earth System Modelling, *Geoscientific Model Development*, 3(1), 87–104, doi:10.5194/gmd-
6 3-87-2010, 2010.

7 Smith, R., Jones, P., Briegleb, B., Bryan, F., Danabasoglu, G., Dennis, J., Dukowicz, J., Eden,
8 C., Fox-Kemper, B., Gent, P., Hecht, M., Jayne, S., Jochum, M., Large, W., Lindsay, K.,
9 Maltrud, M., Norton, N., Peacock, S., Vertenstein, M., and Yeager, S.: The Parallel Ocean
10 Program (POP) reference manual ocean component of the Community Climate System Model
11 (CCSM) and Community Earth System Model (CESM), Los Alamos National Laboratory,
12 LAUR-10-01853, available at
13 <http://www.cesm.ucar.edu/models/cesm1.1/pop2/doc/sci/POPRefManual.pdf> (last access: 15
14 October 2015), 141 p., 2010.

15 Valcke, S., Balaji, V., Craig, A., DeLuca, C., Dunlap, R., Ford, R. W., Jacob, R., Larson, J.,
16 O’Kuinghttons, R., Riley, G. D., and Vertenstein, M.: Coupling technologies for Earth
17 System Modelling, *Geoscientific Model Development*, 5(6), 1589–1596, doi:10.5194/gmd-5-
18 1589-2012, 2012.


19 Valcke, S.: The OASIS3 coupler: a European climate modelling community software,
20 *Geoscientific Model Development*, 6(2), 373–388, doi:10.5194/gmd-6-373-2013, 2013.

21 Valcke, S., Craig, T., and Coquart, L.: The OASIS3-MCT parallel coupler, in: The Second
22 Workshop on Coupling Technologies for Earth System Models (CW2013), available at:
23 https://wiki.cc.gatech.edu/CW2013/images/a/a0/OASIS_MCT_abstract.pdf (last access: 15
24 October 2015), 2013.

25 Valcke, S., Craig, T. and Coquart, L.: OASIS3-MCT User Guide, OASIS3-MCT_3.0,
26 Technical Report TR/CMGC/15/38, Cerfacs, France, 2015
27

Algorithm 1. Generating the parallel decompositions of the sender and the receiver according to an average communication depth of the sender in the P2P implementation.


Input	Number of processes of the sender: M Number of processes of the receiver: N Number of points in the grid: $Grid_pnts$ Average communication depth per process of the sender in the P2P implementation: $Avg_send_depth, Avg_send_depth \leq N$ The flag that specifies whether the communication depths among processes are the same: $Is_balanced$
Output	Parallel decomposition of the sender Parallel decomposition of the receiver
1	Determine the parallel decomposition of the sender Considering that the numbers of grid points are always balanced among the processes of a model, assign around $Grid_pnts/M$ grid points to each process of the sender.
2	Determine the communication depth of each process of the sender
2.1	If the flag $Is_balanced$ is set to true, set the communication depth of each process of the sender to be Avg_send_depth ;
2.2	Otherwise, randomly determine the communication depth of each process of the sender
2.2.1	Initialize the communication depth of each process of the sender to be 1
2.2.2	Randomly select a process of the sender whose communication depth does not exceed N and $Grid_pnts/M$, and then increase its communication depth by 1, until the average communication depth of all processes of the sender reaches Avg_send_depth .
3	Determine the grid points of each communication For each process of the sender, assign the corresponding grid points to all communications of this process (a grid point belongs to only one communication)
3.1	If the flag $Is_balanced$ is set to true, assign the grid points to all communications evenly.
3.2	Otherwise, assign the grid points to each communication randomly
3.2.1	Assign one grid point to each communication
3.2.2	For each of remaining grid points, randomly select a communication for it
4	Determine the parallel decomposition of the receiver through assigning the grid points in each communication to a process of the receiver For each process of the sender, assign the grid points in each communication of it to a distinct receiver process: to make the numbers of grid points balance among the processes of the receiver in the final parallel decomposition, a communication with bigger number of grid points will be assigned to a receiver process with smaller total number of grid points that have been assigned to it.



1

2 Figure 1. Average execution time of the P2P implementation when transferring 14 2-D fields
3 from CLM3 to GAMIL2. In each test, the atmosphere model GAMIL2 and the land surface
4 model CLM3 use the same number of cores; they do not share the same computing nodes.
5 The horizontal grid of the 14 2-D fields contains 7680 (128×60) grid points.


6



1


2 Figure 2. Variation of bandwidth (y-axis) of an MPI P2P communication with respect to the
 3 message size (x-axis). The results are generated from our benchmark. In the benchmark, one
 4 process sends messages with different sizes to the other process. The two processes of the P2P
 5 communication run on two different computing nodes of Tansuo100.

6



1
2
3
4
5


Figure 3. Variation of message size of the P2P implementation (y-axis) in GAMIL2-CLM3 with respect to the number of cores per model (x-axis). The experimental setup is similar to that shown in Fig. 1.



1

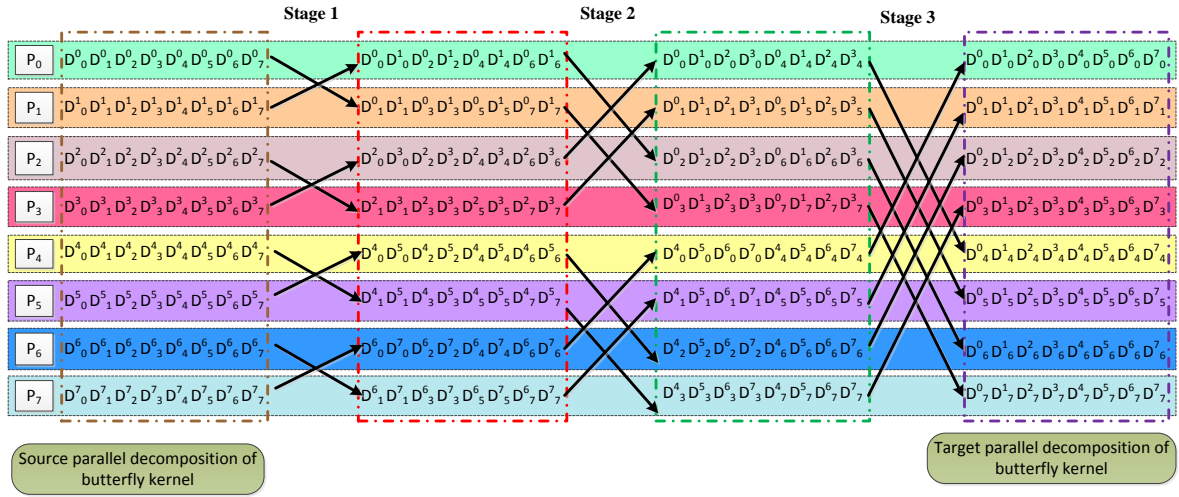
2 Figure 4. Variation of the communication depth of one process (y-axis) using the P2P
 3 implementation in GAMIL2-CLM3 with respect to the number of cores per model (x-axis).
 4 The experimental setup is similar to that shown in Fig. 1.

5




1
 2 Figure 5. Ideal and actual bandwidths of the P2P implementation (y-axis) in GAMIL2-CLM3
 3 when gradually increasing the number of cores used by each component model (x-axis). The
 4 experimental setup is similar to that shown in Fig. 1. The ideal bandwidth is calculated from
 5 the message size and the MPI bandwidth measured in Fig. 2; and the actual bandwidth is
 6 calculated from Fig. 1.

7
 8



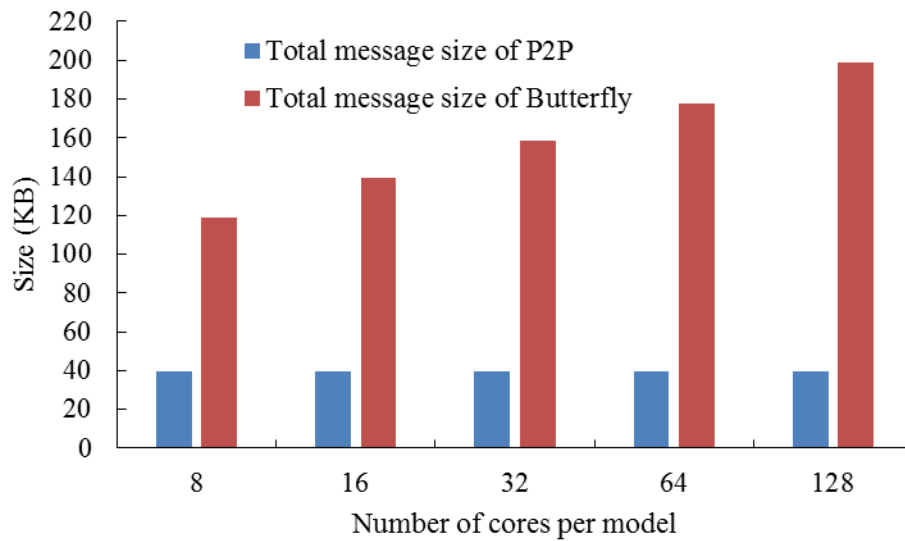
1
 2 Figure 6. An example of the butterfly kernel with eight processes. Each colored row stands
 3 for one process (P_0 - P_7). There are multiple stages (each column of arrows represents a stage
 4 (Stage 1 to Stage 3)) in the butterfly kernel. Each arrow stands for an MPI P2P
 5 communication from one process to another. D_j^i means the data is originally in process P_i
 6 according to the source parallel decomposition and is finally in process P_j according to the
 7 target parallel decomposition.

8




1
2
3
4
5
6

Figure 7. The butterfly implementation, which is composed of three parts: the butterfly kernel; process mapping from the sender to the butterfly kernel; and process mapping from the butterfly kernel to the receiver.




1
 2 Figure 8. Total message size transferred by P2P implementation and butterfly implementation
 3 (y-axis) in GAMIL2-CLM3, when varying the number of cores used by each model (x-axis).
 4 The experimental setup is similar to that shown in Fig. 1.
 5



1
2
3
4
5


Figure 9. Average message size transferred by P2P implementation and butterfly implementation (y-axis) in GAMIL2-CLM3, when varying the number of cores used by each model (x-axis). The experimental setup is similar to that shown in Fig. 1.




1

2 Figure 10. Maximum communication depth, average communication depth and minimum
 3 communication depth in P2P implementation and butterfly implementation (y-axis), when
 4 varying the number of cores used by each model (x-axis) in GAMIL2-CLM3. The
 5 experimental setup is similar to that shown in Fig. 1.


6



(a)



(b)



(c)

1

2 Figure 11. An example of process mappings, given that the sender has five processes (S_0 - S_4),

3 the receiver has 10 processes (R_0 - R_9) (there is no common process between the sender and


4 receiver), and the butterfly kernel contains eight processes (B_0 - B_7). Panels (a) and (b) show

5 how to iteratively pair processes of the sender and receiver, respectively. There are

6 multiple stages in the iterative pairing of processes of the sender and receiver. In each stage,

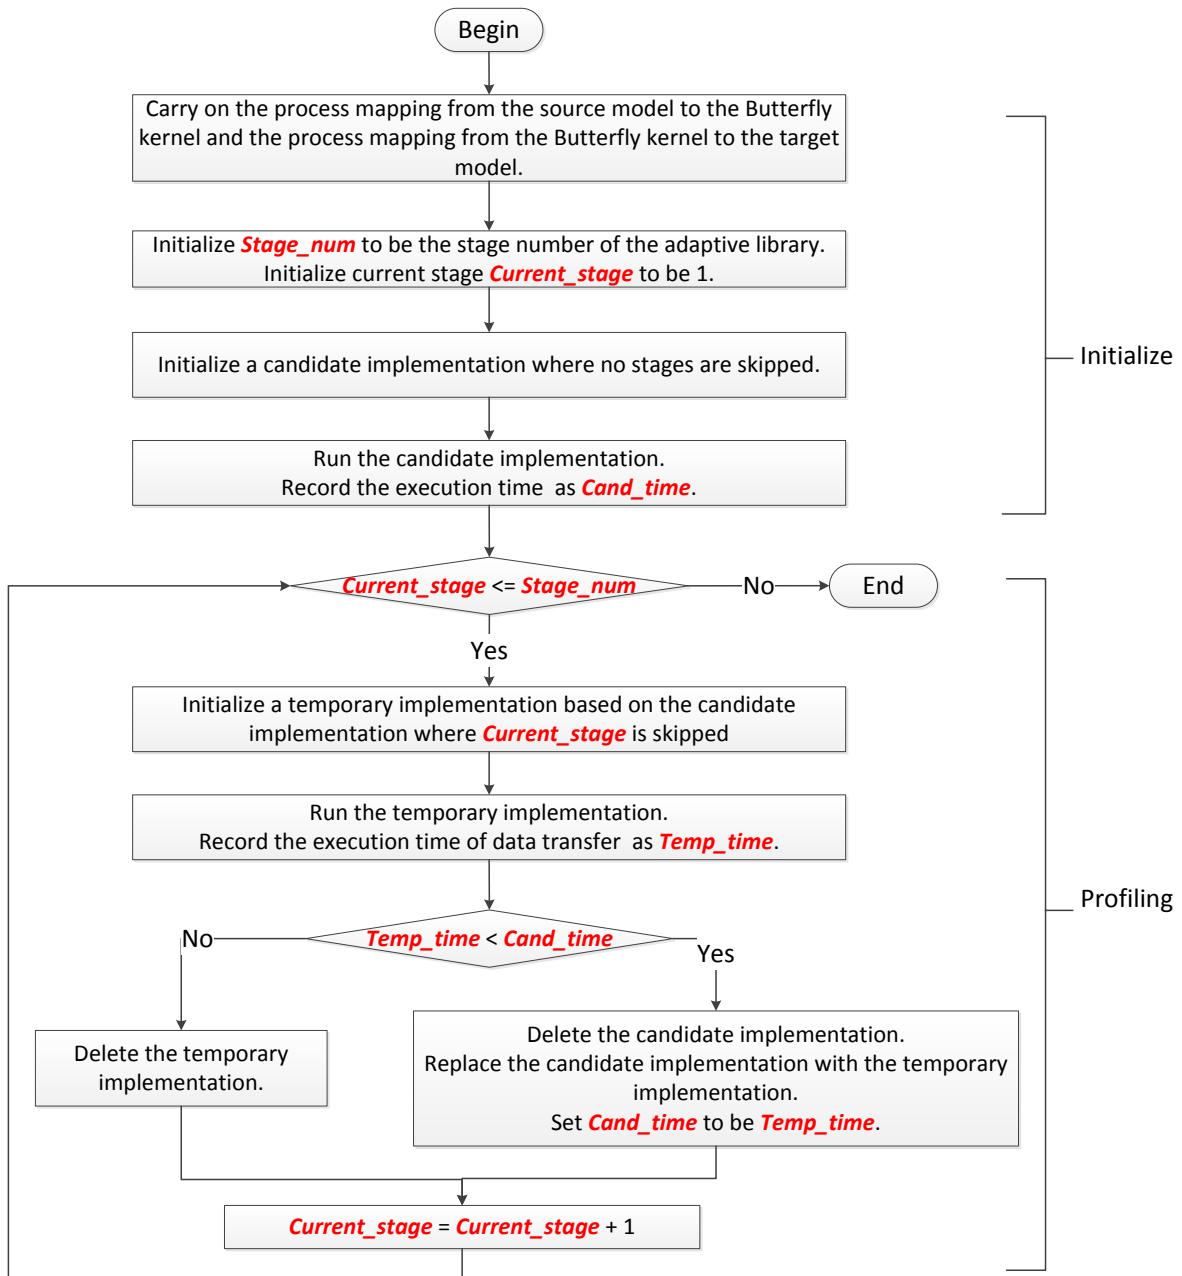
7 the processes in the same color are grouped into one process pair. Panel (c) shows how to map

8 the reordered processes of the sender and receiver onto the processes of the butterfly kernel.



1
2
3
4
5


Figure 12. An example of the adaptive data transfer library with eight processes, where Stage 2 of the butterfly implementation is skipped and replaced by P2P communication of three MPI messages per process.



1


2 Figure 13. A flowchart for determining an appropriate implementation of the adaptive data
 3 transfer library.

4

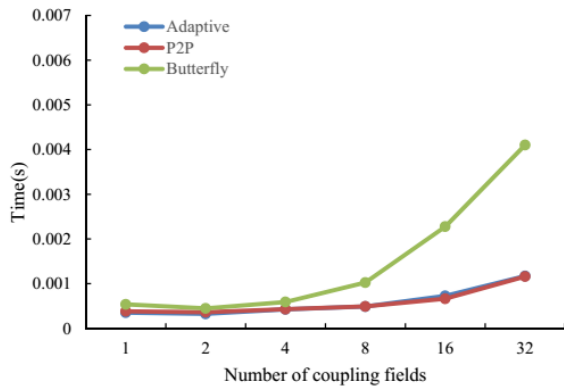


1
2 Figure 14. Initialization time (y-axis) of one data transfer between two toy models using a
3 rectangular grid (of 192×96 grid points) when varying the number of cores used by each toy
4 model (x-axis). There are 10 2-D coupling fields transferred from the source toy model to the
5 target toy model. In each test, all processes of the sender in the P2P implementation have the
6 same communication depth. If the number of cores per model used is less than 24, the
7 communication depth per sender process in the P2P implementation is equal to the number of
8 cores per model; otherwise, the communication depth per sender process in the P2P
9 implementation is 24. The parallel decompositions of the sender and the receiver for a given
10 setting of communication depth are generated by Algorithm 1.

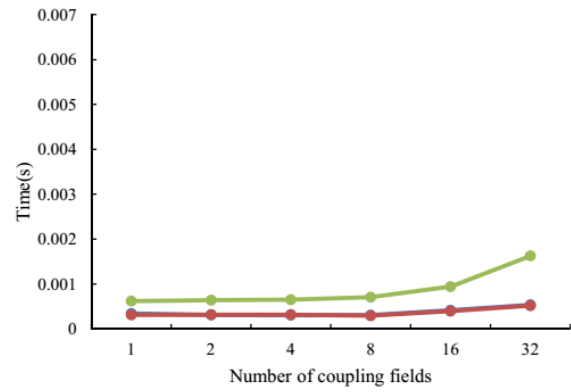
11




1
 2 Figure 15. Average execution time (y-axis) of one data transfer between two toy models with
 3 the same rectangular grid (of 192×480 grid points) when varying the communication depth
 4 per sender process in the P2P implementation (x-axis). Each toy model is run with 1024 cores.
 5 There are 10 2-D coupling fields transferred from the source toy model to the target toy model.
 6




(a)



(b)

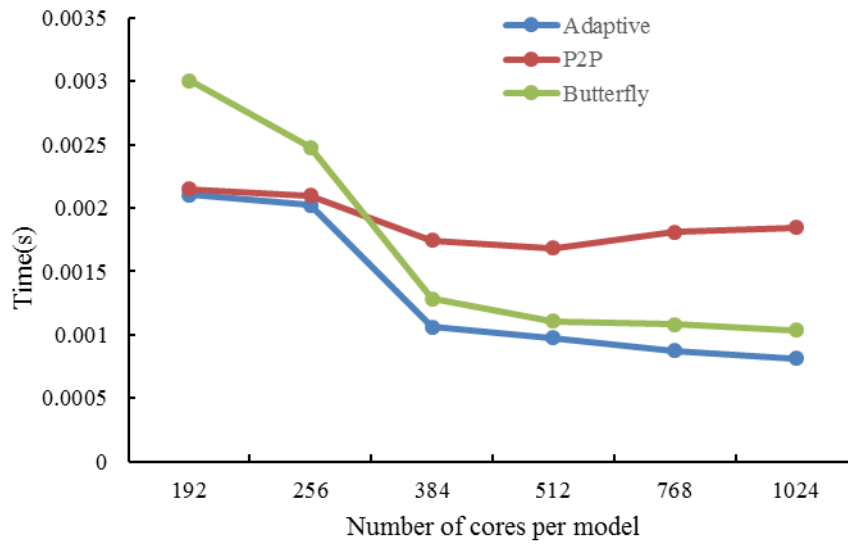


(c)



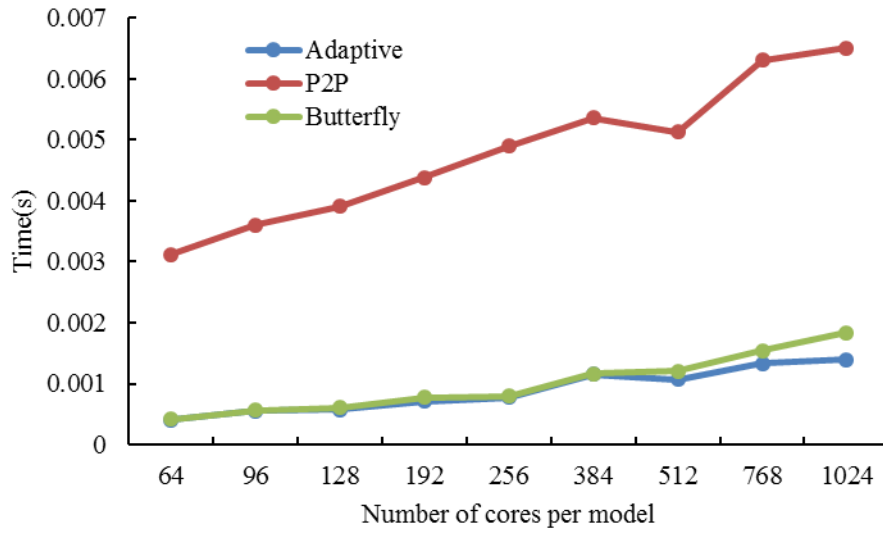
(d)

1
 2 Figure 16. Average execution time (y-axis) of one data transfer between two toy models with
 3 the same rectangular grid (of 192×480 grid points) when varying the number of coupling
 4 fields transferred (x-axis). There are four simulation tests for the evaluation. In simulation (a),
 5 each toy model is run with 256 cores, and the communication depth per sender process in the
 6 P2P implementation is 12. In simulation (b), each toy model is run with 1024 cores, and the
 7 communication depth per sender process is in the P2P implementation 12. In simulation (c),
 8 each toy model is run with 256 cores, and the communication depth per sender process in the
 9 P2P implementation is 48. In simulation (d), each toy model is run with 1024 cores (or
 10 processes), and the communication depth per sender process in the P2P implementation is 48.
 11



1
2
3
4
5
6
7

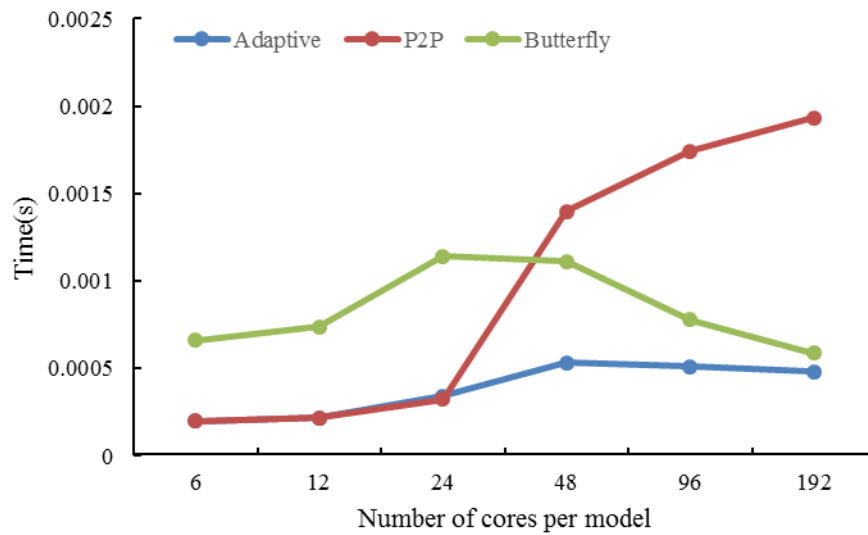
Figure 17. Average execution time (y-axis) of one data transfer between two toy models with the same rectangular grid (of 192×480 grid points) when varying the number of cores used by each toy model (x-axis). There are 10 2-D coupling fields transferred from the source toy model to the target toy model. In each test, the communication depth per sender process in the P2P implementation is 24.



1

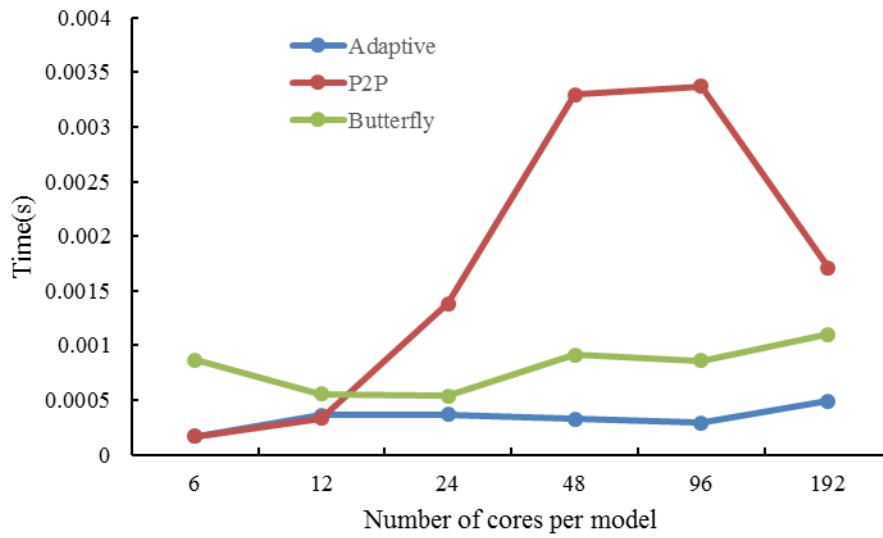
2 Figure 18. Average execution time (y-axis) of one data transfer between two toy models. In
 3 this evaluation, each process (running on a unique processor core) of the toy models have 96
 4 grid points, while different processes have different communication depth and different
 5 message sizes in the P2P implementation. The number of coupling fields transferred is set to
 6 20.

7




1
2
3
4
5
6
7
8

Figure 19. Average execution time (y-axis) of one data transfer between the land surface model CLM4 and the coupler CPL7 in CESM when varying the number of cores used by each model (x-axis): 32 coupling fields on the CLM horizontal grid (the grid size is $144 \times 96 = 13824$) are transferred from the land surface model CLM4 to the coupler CPL7. The performance results of the P2P implementation are obtained through running the adaptive data transfer library when it completely switches to the original P2P implementation.




1
2
3
4
5
6
7

Figure 20. Average execution time (y-axis) of one data transfer between the atmosphere model GAMIL2 and the land surface model CLM3 in GAMIL2-CLM3 when varying the number of cores used by each model (x-axis): 14 coupling fields on the GAMIL2 horizontal grid (the grid size is $128 \times 60 = 7680$) are transferred from the land surface model CLM3 to the atmosphere model GAMIL2.



1
 2 Figure 21. Average execution time (y-axis) of one data rearrangement for the parallel
 3 interpolation from the atmosphere grid (the grid size is $144 \times 96 = 13824$) to the ocean grid (the
 4 grid size is $320 \times 384 = 122880$) in CESM when varying the number of cores used by each
 5 model (x-axis).
 6



1

2 Figure 22. Performance improvement of the coupled model GAMIL2-CLM3 achieved by the
 3 butterfly implementation and the adaptive data transfer library, with the whole model time of
 4 GAMIL2-CLM3 using the P2P implementation as the baseline.