

Toward Adjoinable MPI

Jean Utke* Laurent Hascoët† Chris Hill‡ Paul Hovland§
Uwe Naumann¶

Abstract

Automatic differentiation is the primary means of obtaining analytic derivatives from a numerical model given as a computer program. Therefore, it is an essential productivity tool in numerous computational science and engineering domains. Computing gradients with the adjoint (also called reverse) mode via source transformation is a particularly beneficial but also challenging use of automatic differentiation. To date only ad hoc solutions for adjoint differentiation of MPI programs have been available, forcing automatic differentiation tool users to reason about parallel communication dataflow and dependencies and manually develop adjoint communication code. Using the communication graph as a model we characterize the principal problems of adjoining the most frequently used communication idioms. We propose solutions to cover these idioms and consider the consequences for the MPI implementation, the MPI user and MPI-aware program analysis. The MIT general circulation model serves as a use case to illustrate the viability of our approach.

keywords: MPI, automatic differentiation, source transformation, reverse mode

1 Introduction

In many areas of computational science, it is necessary or desirable to compute the derivatives of functions. In numerical optimization, gradients and sometimes Hessians are used to help locate the extrema of a function. Sensitivity analysis of computer models of physical systems can provide information about how various parameters affect the model and how accurately certain parameters must be measured. Moreover, higher-order derivatives can improve the accuracy of a numerical method, such as a differential equation solver, enabling, for example, longer time steps.

*corresp. author: utke@mcs.anl.gov, University of Chicago and Argonne National Laboratory, Argonne, IL, USA

†INRIA Sophia-Antipolis, Valbonne, France

‡EAPS, MIT, Cambridge, MA, USA

§Mathematics and Computer Science, Argonne National Laboratory, Argonne, IL, USA

¶Department of Computer Science, RWTH Aachen University, Aachen, Germany

31 Automatic differentiation is a technique for computing the analytic deriva-
32 tives of numerical functions given as computer programs. Automatic differ-
33 entiation exploits the associativity of the chain rule and the finite number of
34 intrinsic mathematical functions in a programming language to automate the
35 generation of efficient derivative code [8]. The adjoint (or reverse) mode of au-
36 tomatic differentiation is particularly attractive for computing first derivatives
37 of scalar functions, because it enables one to compute gradients at a cost that
38 is a small multiple of the cost of computing the function and is independent of
39 the number of input variables.

40 The adjoint mode of automatic differentiation combines partial derivatives
41 according to the chain rule, starting at the output (dependent) variable and
42 proceeding (in a direction opposite to the control and data flow of the original
43 function computation) to the input (independent) variables. For any variable
44 u in the original program \mathcal{P} , the automatic differentiation procedure creates
45 an adjoint variable \bar{u} in the adjoint program $\bar{\mathcal{P}}$. This variable represents the
46 derivative of the output variable with respect to u . Consequently, a statement
47 of the form $v=\phi(u)$ in the original program \mathcal{P} results in an update of the form
48 $\bar{u}+=\bar{v}*(\partial v/\partial u)$. Typically ϕ is some intrinsic like `sin`, `cos` etc. in the pro-
49 gramming language of the numerical model to be adjointed. The assignment v
50 $= \phi(u)$ may overwrite a previously used value of v . The generic formulation of
51 the adjoint statement as an increment of the adjoint counterparts of the original
52 right-hand-side arguments necessitates to set $\bar{v}=0$ subsequent to the increment
53 of \bar{u} . Therefore, the simple assignment $v=u$ has as adjoint the two statements
54 $\bar{u}+=\bar{v}; \bar{v}=0$. In the following we will see that this plays an important role in the
55 practical implementation of message-passing adjoints. Because the derivative of
56 v with respect to u , $\partial v/\partial u$, may depend on the value of u and because the vari-
57 able u may be reused and overwritten many times during the function evaluation,
58 the derivative code must record or recompute all overwritten variables whose
59 value is needed in derivative computation. In practice, domain-specific data
60 flow analysis is used to identify variables whose values must be recorded, partial
61 derivatives are “pre-accumulated” within basic blocks, and complex incremental
62 and multilevel checkpointing schemes are employed to reduce memory require-
63 ments [21]. However, for simplicity and without loss of generality, in this paper
64 we assume that a program (or program section) \mathcal{P} is transformed into a new
65 program section $\mathcal{P}^* = \mathcal{P}^+\bar{\mathcal{P}}$, where \mathcal{P}^+ runs forward, recording all overwritten
66 variables, and $\bar{\mathcal{P}}$ runs backward, computing partial derivatives and combining
67 them according to the chain rule. The “backward” execution is accomplished
68 by reversing the flow of control. This implies a reversal of the statement order
69 within basic blocks including calls to communication library subroutines.

70 Many large-scale computational science applications are parallel programs
71 implemented by using MPI message passing. Consequently, in order to apply
72 the adjoint mode of automatic differentiation to these applications, mechanisms
73 are needed that reverse the flow of information through MPI messages. Previous
74 work [2, 3, 9, 12, 13] has examined the automatic differentiation of parallel pro-
75 grams, but this work has focused primarily on the forward mode of automatic
76 differentiation or has relied on the user to implement differentiated versions of

77 communication routines or other ad hoc methods. In this paper we introduce
78 a mechanism for the adjoint mode differentiation of MPI programs, including
79 MPI programs that use nonblocking communication primitives. Given this con-
80 text we focus on qualitative statements regarding the ability to automatically
81 create adjoint code for the most common MPI idioms and the preservation of
82 the basic characteristics of the communication idiom. The latter plays a role
83 in ensuring the correctness of the transformation and retaining the generic per-
84 formance advantages for which a given MPI idiom may have been chosen in
85 the original model. Incremental runtime improvements or suggestions how to
86 improve the communication interface as a whole are beyond the scope of this pa-
87 per. Consequently, we do not present the timings of the forward and the adjoint
88 communication. Instead, we use a case study to show the principal feasibility
89 of our approach.

90 In Sec. 2 we introduce the MPI idioms of concern in this paper. Section 3
91 briefly covers the transformation of plain point-to-point communication and
92 details possible solutions for more complex idioms that are the main contribution
93 of this paper. In Sec. 4 we highlight how the approach was used to automate
94 the transformation of the communication in the MIT general circulation mode.
95 We summarize the results in Sec. 6.

96 2 MPI Idioms and Program Transformation

97 In the following sections we omit the `mpi_` prefix from subroutine and variable
98 names and also omit parameters that are not essential in our context. This
99 section briefly introduces the message-passing concepts relevant to our subject.
100 An automatic transformation of message-passing logic has to be aware of the
101 efficiency considerations that are the reason for different communication modes
102 and the constraints that are implied by these communication modes.

103 Two commonly used models for message-passing communications are the
104 MPI control flow graph (MPICFG) [18] and the communication graph [19, pp.
105 399–403]. A central issue for correct MPI programs is to be deadlock free.
106 Deadlocks can be visualized as cycles in the communication graph. A cycle
107 indicating a deadlock and the use of reordering and buffering to resolve it are
108 shown in Fig. 1. For the plain (frequently called “blocking”) pairs of `send/recv`
109 calls, the edges linking the vertices are bidirectional because the MPI standard
110 allows a blocking implementation; that is, the `send/recv` call may return only
111 after the control flow in the counterpart has reached the respective `recv/send`
112 call. In complicated programs the deadlock-free order may not always be appar-
113 ent. For large data sets the buffered `send` may run out of buffer space, thereby
114 introducing a deadlock caused by memory starvation.

115 A third option to resolve the deadlock shown in Fig. 2(left) uses the nonblock-
116 ing `isend(a,r)`, which keeps the data in the program address space referenced
117 by variable `a` and receives a request identifier `r`. The program can then advance
118 to the subsequent `wait(r)`, after whose return the data in the send buffer `a` is
119 known to have been transmitted to the receiving side.

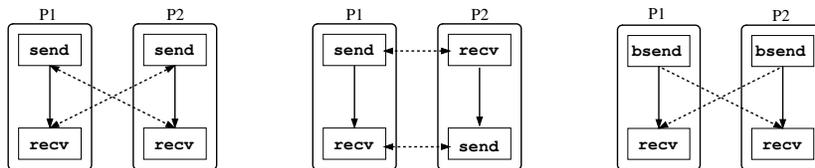


Figure 1: Two processes P1 and P2 want to send data to each other at the same time and deadlock (left). We can reorder the calls (center) to break the deadlock. We could instead keep the order but unblock the send call by using the “buffered” version `bsend`, thus making the communication dependency edges unidirectional (right) and removing the cycle.

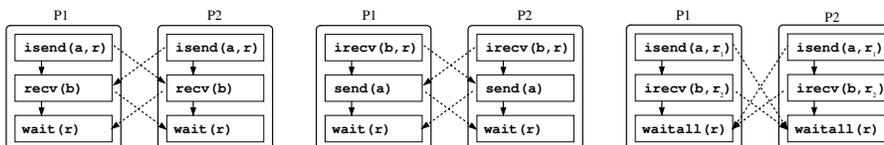


Figure 2: We use the nonblocking send variant `isend` followed by `wait` to break the deadlock scenario in Fig. 1 (left). The same result can be achieved using the nonblocking `irecv` (center) or a combination of the two (right), where we also group the `wait` operations into `waitall`.

120 We assume here that the input program \mathcal{P} is deadlock free. However, the
 121 automatic transformation has to ensure that the transformed MPI program $\bar{\mathcal{P}}$
 122 is also deadlock free. Thus, the transformation has to be cognizant of specific
 123 communication patterns in \mathcal{P} to retain their ability to break potential deadlocks.

124 Other than permitting an immediate return, the variety of different modes
 125 for `send` (and `recv`) calls has its rationale in efficiency considerations. Unlike
 126 `bsend`, an `isend` avoids copying the data to an intermediate buffer but also
 127 requires that the send buffer not be overwritten until the corresponding `wait`
 128 call returns. Similarly, for an `irecv`, see Fig. 2(center), a read (or overwrite) of
 129 the receive buffer prior to the return of the corresponding `wait` returns yields
 130 undefined values. While the transformation should retain efficiency advantages
 131 for $\bar{\mathcal{P}}$, it also has to satisfy the restrictions on the buffers. Because one would
 132 like to minimize artificially imposed order on the message handling, often the
 133 individual `wait` calls are collected in a single `waitall` call; see Fig. 2(right).
 134 The `waitall` vertices in the communication graph typically have more than
 135 one communication in-edge. Two other common scenarios causing multiple
 136 communication in- and out-edges are collective communications (for instance,
 137 broadcasts and reductions) and the use of wildcard for the tag or the source
 138 parameter. In Sec. 3 we explain the consequences of multiple communication
 139 in- and out-edges.

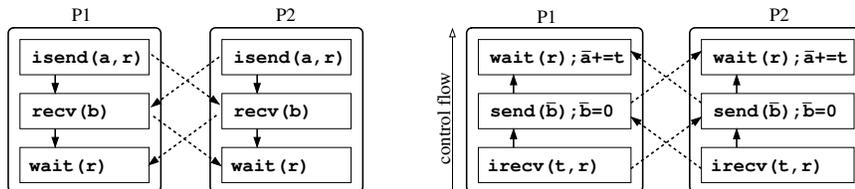


Figure 3: As in Fig. 2 we use `isend` to break a deadlock. This communication pattern can be adjointed while remaining deadlock free by replacing the MPI calls and reversing the direction of the communication edges and the control flow (right).

140 3 Adjoining MPI

141 In this section we explain the construction of the adjoint $\bar{\mathcal{P}}$ of our program
 142 section of interest \mathcal{P} . A direct application of a source transformation tool to
 143 an MPI implementation is impractical for many reasons. One obvious reason is
 144 that we would merely shift the need to prescribe adjoint semantics to commu-
 145 nication operations to some lower level not covered by the MPI standard. For
 146 the transformation we will consider certain patterns of MPI library calls and
 147 propose a set of slightly modified interfaces that we can then treat as atomic
 148 units in a transformation that implements the adjoint semantic.

149 As in sequential programs, the adjoint $\bar{\mathcal{P}}$ will require certain variable values
 150 during its execution. These values might have been recorded in the accompany-
 151 ing augmented forward section \mathcal{P}^+ . However, the particular means of restoring
 152 these values is not the subject of this paper and does not affect what is proposed
 153 here. Consequently, we do not specify \mathcal{P}^+ for the following examples, and we
 154 omit from $\bar{\mathcal{P}}$ any statements related to restoring the values.

155 One can consider a `send(a)` to be a use of the data in variable `a` and the
 156 corresponding `recv(b)` into a variable `b` to be a setting of the data in `b` that
 157 is equivalent to writing a simple assignment statement `b=a`. As explained in
 158 Sec. 1 the adjoint statements corresponding to this assignment are `$\bar{a}+=\bar{b}$` ; `$\bar{b}=0$` .
 159 Applying the above analogy we can express the adjoint as `$\text{send}(\bar{b})$` ; `$\bar{b}=0$` as
 160 the adjoint of the original `recv` call and `$\bar{a}+=\bar{t}$` as the adjoint of the
 161 original `send` call. This has been repeatedly discovered and used in various
 162 contexts (e.g., [3, 4]) and is the extent to which automatic transformation has
 163 been supporting MPI until now.

164 3.1 Required Context

165 The semantics of the adjoint computation as introduced in Sec. 1 implies that
 166 both the control flow and the communication edges have to be reversed. We
 167 already mentioned the need to preserve certain features of the communication
 168 patterns to keep the communication efficient and deadlock free; see, for exam-
 169 ple, Fig. 3. Considering the modes of `send` and `recv` calls, we can derive a

Table 1: Rules for adjoining a restricted set of MPI `send/recv` patterns. We omit all parameters except for the buffers `a`, `b`, a temporary buffer `t`, and the request parameter `r` for nonblocking calls.

	in \mathcal{P}		in $\bar{\mathcal{P}}$	
	call	paired with	call	paired with
1	<code>isend(a,r)</code>	<code>wait(r)</code>	<code>wait(r);a+=t</code>	<code>irecv(t,r)</code>
2	<code>wait(r)</code>	<code>isend(a,r)</code>	<code>irecv(t,r)</code>	<code>wait(r)</code>
3	<code>irecv(b,r)</code>	<code>wait(r)</code>	<code>wait(r);b=0</code>	<code>isend(b,r)</code>
4	<code>wait(r)</code>	<code>irecv(b,r)</code>	<code>isend(b,r)</code>	<code>wait(r)</code>
5	<code>bSEND(a)</code>	<code>recv(b)</code>	<code>recv(t);a+=t</code>	<code>bSEND(b)</code>
6	<code>recv(b)</code>	<code>bSEND(a)</code>	<code>bSEND(b);b=0</code>	<code>recv(t)</code>
7	<code>ssend(a)</code>	<code>recv(b)</code>	<code>recv(t);a+=t</code>	<code>ssend(b)</code>
8	<code>recv(b)</code>	<code>ssend(a)</code>	<code>ssend(b);b=0</code>	<code>recv(t)</code>

170 set of patterns where simple rules suffice for the adjoint program transforma-
171 tion. For simplicity we consider `send(a)`; to be equivalent to `isend(a,r)`;
172 `wait(r)`; and similarly for `recv`. For `send(a)`, that is `isend(a,r)`; `wait(r)`,
173 we apply rule 1 and reverse the control flow, obtaining `irecv(t,r)`; `wait(r)`;
174 `a+=t`, that is, `recv(t)`; `a+=t`. When all communication patterns in a program
175 \mathcal{P} match one of the rules listed in Table 1, then one can replace the respective
176 MPI calls as prescribed. Together with control flow reversal orchestrated by
177 the regular adjoint transformation, the correct reversal of the communication
178 edges then is implied. A framework to formally prove these rules can be found
179 in [17] and we prove a specific case in Sec. 3.4. As evident from the table entries
180 the proper adjoint for a given call depends on the context in the original code.
181 One has to facilitate the proper pairing of the `isend/irecv` calls with their
182 respective individual `waits` for rules 1–4 (intra-process) and also of `send` mode
183 for a given `recv` for rules 5–8 (inter-process). An automatic code analysis will
184 often be unable to determine the exact pairs. Instead one could either use the
185 notion of communication channels identified by pragmas [12] or wrap the MPI
186 calls into a separate layer. This layer needs to encapsulate the required context
187 information (e.g., via distinct `wait` variants) and potentially passes the respec-
188 tive user space buffer as an additional argument; for example, `sendwait(r,a)`
189 may be paired with `isend(a,r)`. Likewise the layer would introduce distinct
190 `recv` variants; for example, `brecv` would be paired with `bSEND`. Note that combi-
191 nations of nonblocking, synchronous and buffered send and receive modes not
192 listed in Table 1 can be easily derived. For instance, the adjoint of a sequence
193 of `ibSEND(a,r)` - `recv(b)` - `wait(r)` involves rule 2 for the `wait` and rule 5
194 for the `recv`, resulting in the adjoint sequence `irecv(t,r)` - `bSEND(b);b=0` -
195 `wait(r);a+=t`.

196 **3.2 Wildcards and Collective Communication**

197 The adjoining recipes have so far considered only cases where the vertices in
 198 the communication graph have single in- and out-edges. Using the MPI wild-
 199 card values for parameters `source` or `tag` implies that a given `recv` might be
 200 paired with any `send` from a particular set; that is, the `recv` vertex has multiple
 201 communication in-edges only one of which at any time during the execution is
 202 actually traversed. Transforming the `recv` into a `send` for the adjoint means
 203 that we need to be able to determine the destination. A simple solution is to
 204 *record* the values of the tag and source parameters in the augmented forward
 205 version \mathcal{P}^+ at run-time.¹ Conceptually this could be interpreted as a run-time
 206 incarnation of the communication graph in which the set of potential in- or out-
 207 edges has been replaced by the one communication that actually takes place.
 208 Thus, the single in- and out-edge property is satisfied again. In \mathcal{P} the wildcard
 209 parameters are replaced with the actual values that were recorded during the
 210 execution of \mathcal{P}^+ , thus ensuring that we traverse the correct, inverted communi-
 211 cation edge. One can show that for any deadlock-free run-time incarnation of
 212 the communication graph, one can construct a corresponding adjoint commu-
 213 nication graph that will also be deadlock free.

214 For collective communications the transformation of the respective MPI calls
 215 is essentially uniform across the participating calls. To illustrate the scenario,
 216 we can consider a summation reduction followed by a broadcast of the result,
 217 which could be accomplished by calling `allreduce` but here we want to do it
 218 explicitly. In \mathcal{P} we sum up to the rank 0 process `reduce(a, b, +)` (i.e. $\mathbf{b}_0 = \sum \mathbf{a}_i$)
 219 followed by `bcast(b)` (i.e. $\mathbf{b}_i = \mathbf{b}_0 \forall i$). The corresponding adjoint statements in
 220 $\bar{\mathcal{P}}$ with a temporary variable `t` and reversed control flow are `t0 = $\sum \bar{\mathbf{b}}_i$` followed
 221 by `$\bar{\mathbf{a}}_i += \mathbf{t}_0 \forall i$` , which, expressed as MPI calls, are `reduce($\bar{\mathbf{b}}$, \mathbf{t} , +)` followed by
 222 `bcast(t); $\bar{\mathbf{a}}_i += \mathbf{t}$` . In short, a reduction becomes a broadcast and vice versa. The
 223 respective communication graphs are shown in Fig. 4. For the purpose of the
 224 automatic differentiation program transformation the multiple communication
 225 in- or out-edges can conceptually be treated as single communication channel
 226 visualized as a hyperedge in the communication graph.²

227 To expose an efficiency concern, we modify the above example slightly to
 228 perform a product reduction instead of the summation. The transformation
 229 remains the same except for the increments `$\bar{\mathbf{a}}_i += (\partial \mathbf{b}_0 / \partial \mathbf{a}_i) \mathbf{t}_0 \forall i$` that follow the
 230 `bcast` in $\bar{\mathcal{P}}$. The above formula for the $\bar{\mathbf{a}}_i$ does not suggest how exactly to
 231 compute the partials $\partial \mathbf{b}_0 / \partial \mathbf{a}_i$. In principle, the partials could be explicitly
 232 computed by using prefix and suffix reduction operations during the recording
 233 sweep [13]. Alternatively one could record the \mathbf{a}_i per process in \mathcal{P}^+ and then
 234 in $\bar{\mathcal{P}}$ first restore the \mathbf{a}_i , then compute all the intermediate products from the
 235 leaves to the root in the reduction tree, followed by propagating the adjoints
 236 from the root to the leaves [1]. This approach requires only two passes over the
 237 tree and thus is less costly than any approach using the explicit computation

¹ They are returned with the `status` of the `recv` call.

² Assuming a hyperedge in general might suggest a synchronization among the participating collective calls, which is not required by the standard.

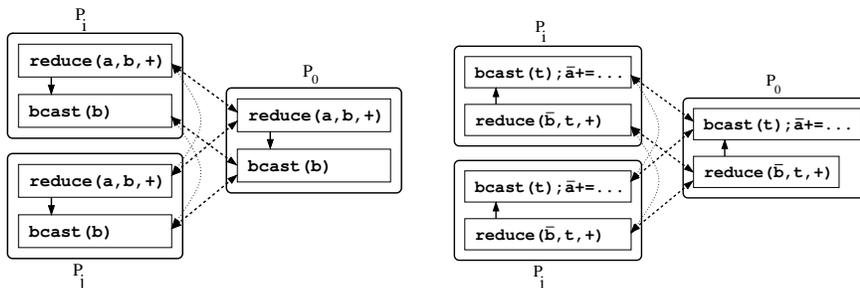


Figure 4: Collective operations reduction and broadcast in \mathcal{P} are all connected with communication edges among themselves (left). If we can identify the rank 0 process among the collective calls, the dotted communication edges can be removed. The adjoint inverts the control flow but keeps the bidirectional communication edges grouped in the same hyper edges in $\bar{\mathcal{P}}$ (right).

238 of the partials. Unlike the explicit partials computation using pre- and postfix
 239 reductions, MPI does not provide interfaces facilitating the two-pass approach;
 240 consequently, one would have to implement it from scratch.

241 3.3 Grouping wait Operations

242 The grouping of sets of `wait` operations into a call to `waitall` or `waitsome`
 243 can increase the communication efficiency by removing the often artificial or-
 244 der among the requests. It also leads to multiple communication in-edges; see
 245 Fig. 5(left).³ Typically, more than one or even all of these in-edges are traversed,
 246 thereby distinguishing the scenario from the wildcard receive case we considered
 247 in Sec. 3.2. A transformation solely based on the rules in Table 1 would require
 248 first modifying \mathcal{P} such that all the grouped `wait` operations are split into in-
 249 dividual `waits`. While they could then be transformed into the respective
 250 `isend` and `irecv` calls shown in Fig. 5(right), we would in the process lose the
 251 potential performance advantage that prompted the use of `waitall` in the first
 252 place. Without loss of generality we consider a sequence of `isend` calls, followed
 253 by a sequence of `irecv` calls, followed by a `waitall`, as shown in Fig. 5. While
 254 there are no communication edges directly between the `isends` and `irecvs`, we
 255 know that in principle we want to turn send into receive operations and vice
 256 versa. Replacing `isends` and `irecvs` in \mathcal{P} with `irecvs` and `isends` in $\bar{\mathcal{P}}$ begs
 257 the question of where in $\bar{\mathcal{P}}$ the corresponding `waits` should go. This gives us
 258 the rationale to introduce a symmetric counterpart to `waitall` into \mathcal{P} that
 259 we call `awaitall`, which stands for *anti-waitall*. We illustrate the scenario in
 260 Fig. 6(left). In \mathcal{P} there no semantics are assigned to `awaitall`, and the vertex
 261 and the communication edges can be considered nonoperational. The adjoint
 262 transformation shown in Fig. 6(right) makes them operational and in symmet-

³ For simplicity we assume all processes have the same behavior, and we show only the condensed MPI CFG [18].

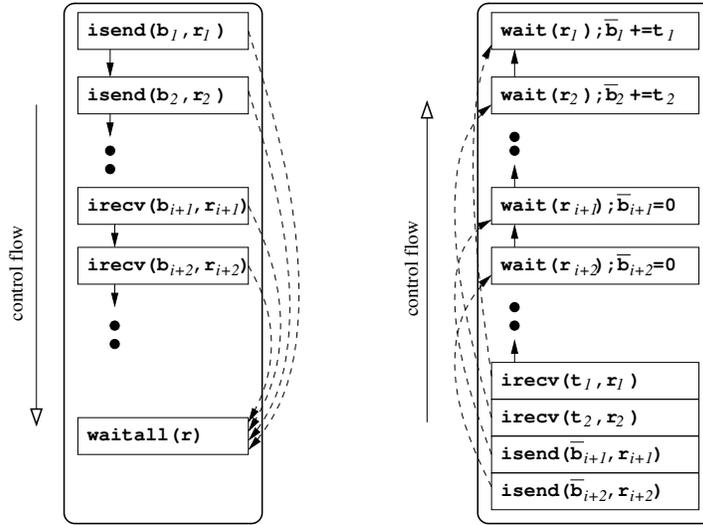


Figure 5: The completion `waitall` has multiple communication in-edges that make a simple vertex based adjoint transformation impossible (left). Without further information, the `waitall` is split into individual waits, producing a less efficient adjoint code (right) where there is no `waitall`.

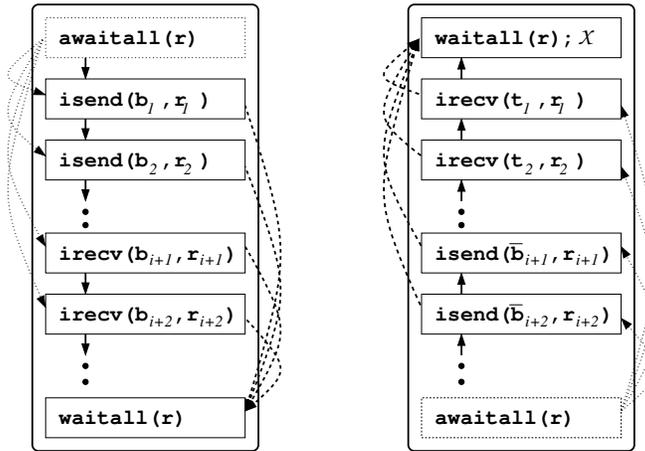


Figure 6: We introduce into \mathcal{P} a symmetric nonoperational counterpart to `waitall` called `awaitall` (left). The adjoint transformation to $\bar{\mathcal{P}}$ turns `awaitall` into `waitall`, `isend` into `irecv` and vice versa (right).

263 rical fashion renders nonoperational the out-edges of the `awaitall` vertex in $\bar{\mathcal{P}}$
 264 that corresponds to the `waitall` in \mathcal{P} . The final \mathcal{X} at the top Fig. 6(right)
 265 denotes the adjoint buffer updates $\bar{b}_j += t_j, j = 1, \dots, i$ and $\bar{b}_j = 0, j = i + 1, \dots$
 266 that have to wait for completion of the nonblocking calls. The following section
 267 proves the correctness of this transformation and gives the rationale for
 268 symmetrically extending the restrictions on writing and reading the `isend` and
 269 `irecv` buffers to the entire section between the `awaitall` and the `waitall`.

270 3.4 Correctness of Adjoint Communication Patterns

271 The technique used in this section can be applied to prove all the adjoint trans-
 272 formation recipes discussed in this paper. We illustrate it using one of the
 273 recipes. Consider the *partitioned global address space* (PGAS) versions $\mathcal{P}^=$ of
 274 the message-passing program \mathcal{P} involving n processes p_1, \dots, p_n ; see also [5].
 275 The partitioning augments all program variables with an additional dimension
 276 of length n . Communications are translated into assignments between the aug-
 277 mented program variables. Additional variables are introduced for buffered
 278 communication. Barriers in asynchronous communication yield a set of se-
 279 quentialized versions for a given message-passing program. For example, the
 280 processes p_1 and p_2 performing a single nonblocking send/receive

281 `s0; if(this is p1) isend(a,r); s1; if(this is p2) irecv(b,r); s2; wait(r); s3`

282 with some unspecified statements s_0, \dots, s_3 yield the following six distinct PGAS
 283 variants.

284 **1:** $s_0; s_1^2; b^2 = a^1; s_1^1; s_2; s_3$ **2:** $s_0; s_1; b^2 = a^1; s_2; s_3$
3: $s_0; s_1; s_2^2; b^2 = a^1; s_2^1; s_3$ **4:** $s_0; s_1; s_2^1; b^2 = a^1; s_2^2; s_3$
5: $s_0; s_1^2; s_2^2; b^2 = a^1; s_1^1; s_2^1; s_3$ **6:** $s_0; s_1; s_2; b^2 = a^1; s_3$

285 The superscripts indicate the executing process or address subspace, respec-
 286 tively. Note that $(s_i^1; s_i^2) = (s_i^2; s_i^1)$ as a result of the disjoint address spaces.
 287 Hence, the sequence of statements $s_i; s_{i+1}$ yields the following six semantically
 288 equivalent PGAS sequences:

289 **1:** $s_i^1; s_{i+1}^1; s_i^2; s_{i+1}^2$ **2:** $s_i^2; s_{i+1}^2; s_i^1; s_{i+1}^1$ **3:** $s_i^1; s_i^2; s_{i+1}^1; s_{i+1}^2$
4: $s_i^1; s_i^2; s_{i+1}^2; s_{i+1}^1$ **5:** $s_i^2; s_i^1; s_{i+1}^1; s_{i+1}^2$ **6:** $s_i^2; s_i^1; s_{i+1}^2; s_{i+1}^1$

290 The partial order of the statements is induced by $s_i^j < s_{i+1}^j$. Any two statements
 291 from s_i^j and s_{i+1}^k can be executed in arbitrary order for $j \neq k$. Further combina-
 292 tions resulting from feasible (with respect to data dependence) switches of the
 293 buffer assignment and statements in certain s_i^j lead to an exponential number
 294 of possible actual execution orders that need to be taken into account when
 295 proving properties of PGAS programs. The `isend` restrictions imply that a^1
 296 not be written by s_1 or s_2 . Similarly, b^2 may not be read by s_2 . To prove the
 297 correctness of an adjoint of a message-passing program we need to show that its
 298 adjoint PGAS versions $\overline{(\mathcal{P}^=)}$ are equivalent to the PGAS versions of its adjoint

299 $(\bar{\mathcal{P}})^=$. We will do so for the pattern shown in Fig. 6 rewritten as follows.

$$s_{i-1}; \text{awaitall}(r); s_{i+1}; \dots s_{j-1}; \text{isend}(a, r_\nu); s_{j+1} \dots$$

$$\dots s_{k-1}; \text{irecv}(b, r_\nu); s_{k+1}; \dots s_{l-1}; \text{waitall}(r); s_{l+1}$$

300 We consider n processes for this pattern in \mathcal{P} with $\nu = 1, \dots, n$, $a \neq b$, $j = j(\nu)$,
 301 and $k = k(\nu)$. Reversing the control flow and then applying the recipe yields
 302 for $\bar{\mathcal{P}}$

$$\overline{s_{l+1}}; \text{awaitall}(r); \overline{s_{l-1}}; \dots \overline{s_{k+1}}; \text{Xisend}(\bar{b}, r_\nu); \overline{s_{k-1}} \dots$$

$$\dots \overline{s_{j+1}}; \text{Xirecv}(\bar{a}, r_\nu); \overline{s_{j-1}}; \dots \overline{s_{i+1}}; \text{waitall}(r); \overline{s_{i-1}}$$

where we fold some buffer operations⁴ into the respective send and receive call such that

$$\text{Xisend}(\bar{b}, r_\nu) \equiv \text{isend}(\bar{b}, r_\nu) \rightsquigarrow \bar{b}=0$$

and

$$\text{Xirecv}(\bar{a}, r_\nu) \equiv \text{irecv}(t, r_\nu) \rightsquigarrow \bar{a}+=t$$

303 When we require that a not be written and b not be read by any statement
 304 s_{i+1}, \dots, s_{l-1} , we can show that the $\bar{\mathcal{P}}$ shown above is correct as follows.

305 **Proof:** Consider the possible $\mathcal{P}^=$ versions with placeholders \triangleright and \triangleleft for the
 306 respective wait statements in positions i and l and ∇ for the **irecv** statement
 307 in position k :

$$1: s_{i-1}; \triangleright; s_{i+1} \dots s_{k-1}; \nabla; b^{\nu_2} = a^{\nu_1}; s_{k+1} \dots s_{l-1}; \triangleleft; s_{l+1}$$

$$\vdots$$

$$l-k: s_{i-1}; \triangleright; s_{i+1} \dots s_{k-1}; \nabla; s_{k+1} \dots s_{l-1}; b^{\nu_2} = a^{\nu_1}; \triangleleft; s_{l+1}$$

309 With the reversal of control flow their adjoints are as follows.

$$1: \overline{s_{l+1}}; \triangleleft; \overline{s_{l-1}} \dots \overline{s_{k+1}}; \bar{a}^{\nu_1} += \bar{b}^{\nu_2}; \bar{b}^{\nu_2} = 0; \nabla; \overline{s_{k-1}} \dots \overline{s_{i+1}}; \triangleright; \overline{s_{i-1}}$$

$$\vdots$$

$$l-k: \overline{s_{l+1}}; \triangleleft; \bar{a}^{\nu_1} += \bar{b}^{\nu_2}; \bar{b}^{\nu_2} = 0; \overline{s_{l-1}} \dots \overline{s_{k+1}}; \nabla; \overline{s_{k-1}} \dots \overline{s_{i+1}}; \triangleright; \overline{s_{i-1}}$$

311 Because \mathcal{P} satisfies the **isend/irecv** restrictions, the variable a^{ν_1} is not writ-
 312 ten by any of the statements s_{j+1}, \dots, s_{l-1} , nor is b^{ν_2} read or written by
 313 s_{k+1}, \dots, s_{l-1} . The variable a^{ν_1} may be read by any of the statements
 314 s_{j+1}, \dots, s_{l-1} implying that \bar{a}^{ν_1} may be incremented in $\overline{s_{l-1}}, \dots, \overline{s_{j+1}}$ while \bar{b}^{ν_2}
 315 remains unchanged by $\overline{s_{l-1}}, \dots, \overline{s_{k+1}}$. The only read operation on \bar{a}^{ν_1} within
 316 $\overline{s_{l-1}}, \dots, \overline{s_{j+1}}$ is performed as part of these possibly present increment opera-
 317 tions. In exact arithmetic the increment order for \bar{a}^{ν_1} has no impact on its final
 318 value. Hence, the statement $\bar{a}^{\nu_1} += \bar{b}^{\nu_2}$ followed by $\bar{b}^{\nu_2} = 0$ can be inserted at any
 319 position between $\overline{s_{l+1}}$ and $\overline{s_{k-1}}$. In short, we constructed $(\bar{\mathcal{P}})^=$ and found the
 320 adjoints of all PGAS versions are equivalent.

321 We now consider $(\bar{\mathcal{P}})^=$ where \triangle is the placeholder for the **isend** statement in
 322 position j

⁴ In an implementation, these buffer operations have to be delayed until the data is transmitted or received, respectively.

323

1: $\overline{s_{l+1}}$; \triangleleft ; $\overline{s_{l-1}} \dots \overline{s_{j+1}}$; \triangle ; $t=\bar{b}^{\nu_2}$; $\bar{b}^{\nu_2}=0$; $\bar{a}^{\nu_1}+=t$; $\overline{s_{j-1}} \dots \overline{s_{i+1}}$; \triangleleft ; $\overline{s_{i-1}}$

324

\vdots

$l-k$: $\overline{s_{l+1}}$; \triangleleft ; $\overline{s_{l-1}} \dots \overline{s_{j+1}}$; \triangle ; $\overline{s_{j-1}} \dots \overline{s_{i+1}}$; $t=\bar{b}^{\nu_2}$; $\bar{b}^{\nu_2}=0$; $\bar{a}^{\nu_1}+=t$; \triangleleft ; $\overline{s_{i-1}}$

325

326

327

328

329

330

331

332

333

334

At this point the need to impose additional restrictions on \bar{a}^{ν_1} and \bar{b}^{ν_2} becomes apparent. The temporary variable t is neither read nor written by any of the statements $\overline{s_{l-1}}, \dots, \overline{s_{i-1}}$. It can be removed as the result of copy-propagation. In the $(\bar{\mathcal{P}})^=$ versions the adjoint computation takes place between $\overline{s_{l+1}}$ and $\overline{s_{k-1}}$. The same computation is performed in $(\bar{\mathcal{P}}^=)$ between $\overline{s_{j+1}}$ and $\overline{s_{i-1}}$. Hence, the variable \bar{b}^{ν_2} must not be written by any of the statements in $\overline{s_{j-1}}, \dots, \overline{s_{i+1}}$. Moreover, \bar{a}^{ν_1} may only be incremented by these statements. Therefore, in $\bar{\mathcal{P}}$ the variable a^{ν_1} may not be written and b^{ν_2} may not be read or written by s_{i+1}, \dots, s_{j-1} . This reflects exactly our additional requirement, and therefore the proposed transformation is shown to be correct ■

335

3.5 Placement Flexibility and Implementation Choices

336

337

338

339

340

341

342

343

344

345

346

347

348

349

350

351

352

353

354

355

356

In the program section between the `awaitall` and the `waitall` our augmented restriction on the `isend` and `irecv` buffers is symmetric. Just like the `waitall`, the placement of the `awaitall` will have to be done by the MPI programmer who wishes to use the automatic differentiation transformation. While the main goal of this paper is a transformation via recipes applied at the level of a subroutine call, we note that the restrictions in turn permit some flexibility to move the `isend` and `irecv` calls within this program section. The situation is illustrated in Fig. 7. This would afford the maximal time the message-passing system can spend to process communication requests before further computation in the participating processes is halted pending the return of the respective `waitall`. Unlike the transformation recipes we proposed so far, such a modification of the original program requires detailed data dependency information. With a few exceptions, practical message-passing programs will likely not be amenable to an automatic program analysis that can provide the data dependencies with sufficient accuracy. On the other hand, it is perfectly reasonable to consider as a starting point that communication channels in the program are identified by pragmas. Together with pragmas that serve as semantic placeholders for the `awaitall` position, all the information required to apply the adjoint transformation recipes would be present. Obviously, standard program analyses still are required to establish data dependencies for all the other parameters in the MPI calls, their position in the control flow graph, and so forth.

357

358

359

360

361

362

363

From the above it is obvious that an approach based on pragma-identified communication channels would be the most beneficial for general purposes. It is clearly also a rather complicated choice for something that can also be implemented with a set of subroutines that wrap the MPI calls and supply all the required context information via parameters and context-specific versions. An example, mentioned in Sec. 3.1, is a specific `sendwait`, which takes as an additional parameter the corresponding `isend` call's `sendbuffer`. The wrapper

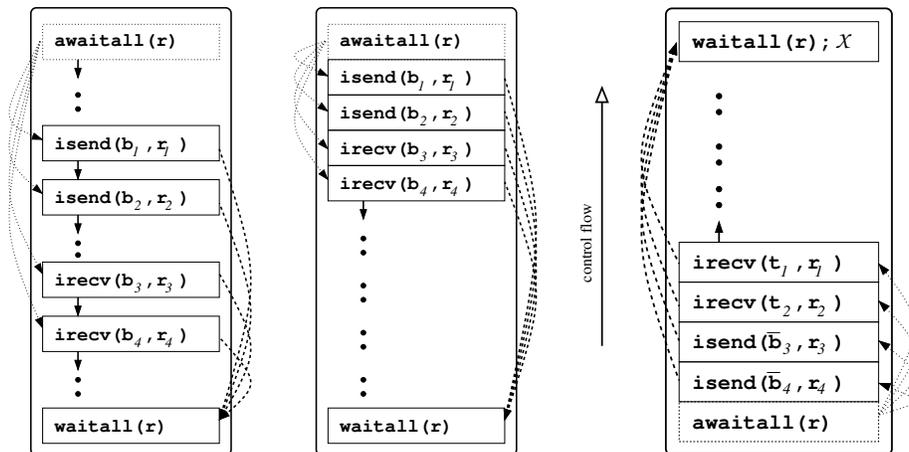


Figure 7: Because the restrictions on the communication buffers are symmetric in \mathcal{P} (left) we can maximize the time between posting a communication request and waiting for the completion by shifting the calls for \mathcal{P}^+ (center) and $\bar{\mathcal{P}}$ (right).

364 routines then can switch their behavior, perhaps via some global setting, between the original and the respective adjoint semantics indicated by the recipe.
 365 Some additional bookkeeping for the delayed buffer operations to be executed by the adjoint semantic for `awaitall` is in principle all that is needed to accomplish the task. More details on this can be found in Sec. 4 related to our
 366 wrapper-based prototype implementation.
 367

370 Motivated by efficiency considerations for computing the partial derivatives in Sec. 3.2 and folding the buffer semantics into an `Xisend` or an `Xirecv` call as in
 371 Sec. 3.4, one might attempt to integrate such interfaces into an MPI implementation, thereby avoiding an extra layer of calls and replicated bookkeeping logic.
 372 Such implementation details are beyond the scope of this paper and subject to ongoing research and discussion in the automatic differentiation community.
 373
 374
 375

376 4 Case Study: MITgcm

377 The MIT General Circulation Model (MITgcm) is an ocean and atmosphere
 378 geophysical fluids simulation code [15, 16] that is widely used for both realistic
 379 and idealized studies and runs on both serial desktop systems and large-scale
 380 parallel systems. It employs a grid-point based, time-stepping algorithm that
 381 derives its parallelism from spatial domain decomposition in two horizontal di-
 382 mensions. Coherence between decomposed regions, in both forward and reverse
 383 mode adjoint computations, is handled explicitly by a set of hand-written com-
 384 munication and synchronization modules that copy data between regions using
 385 either shared-memory or MPI-based messaging. The MITgcm code supports ar-

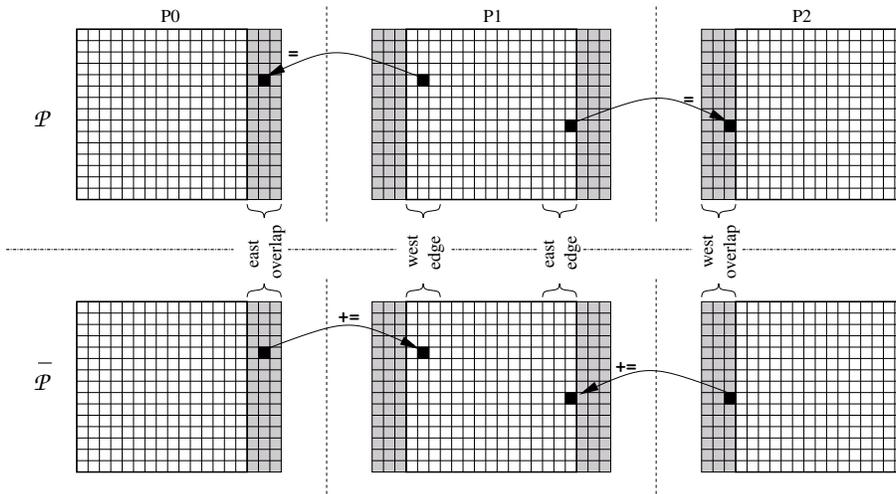


Figure 8: In \mathcal{P} the data at the eastern and western edge of P1 are copied to the respective overlap arrays (shaded gray) of its neighbors P0 and P2. In $\bar{\mathcal{P}}$ the adjoint of this operation is to increment the adjoint edge data in P0 by the adjoint overlap data from its neighbors P0 and P2.

386 arbitrary cost functions [7, 11, 14, 20] for which adjoints can be generated with the
 387 automatic differentiation tools TAF and OpenAD/F. Until now, however, the
 388 automatic adjoint transformation did not extend to the MPI communication
 389 layer. Instead, hand-written “adjoint forms” of the MITgcm communication
 390 and synchronization modules have to be maintained [9, 10] and substituted into
 391 the code generated automatically for the other parts of the MITgcm. The lack
 392 of tool support required other ocean model developers to adopt the same strat-
 393 egy [4]. Creating and maintaining these hand-written adjoint sections are ardu-
 394 ous and highly error-prone tasks, particularly when multiple discretization and
 395 decomposition options require many variants of the communication logic. This
 396 situation provided the impetus to investigate to what extent automatic trans-
 397 formation might support communication patterns that are more sophisticated
 398 than plain `send-recv` pairs.

399 Figure 8 illustrates the approach for a hypothetical two-dimensional array
 400 (the hatched blocks on the diagram) that has been block-decomposed along one
 401 dimension over three processes. The gray-shaded areas, hold copies of data that
 402 is updated by the respective neighbor. In order to maintain coherence, in for-
 403 ward simulation mode, data is sent from the appropriate index regions on P1
 404 (labeled “west edge” and “east edge”) and received by processes P0 and P2
 405 into the index regions (labeled “east overlap” and “west overlap”). The arrows
 406 in the \mathcal{P} row in Fig. 8 illustrate this data traffic for the updates that P1 sends
 407 to its neighbors. The respective adjoint operation is shown in $\bar{\mathcal{P}}$ row. The com-
 408 munication graph for the data exchange in Fig. 8 is shown in Fig. 9. In practice

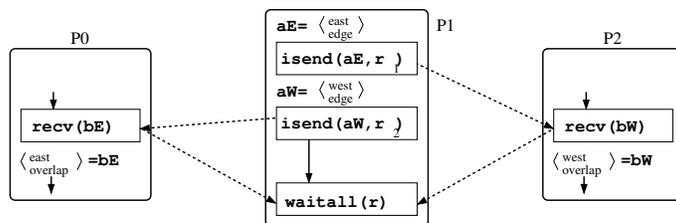


Figure 9: The communication graph for the data exchange in \mathcal{P} illustrated in Fig. 8.

409 the data exchange is of course symmetric (P1 also receives data from P0 and
 410 P2), periodic (P0 and P2 are also neighbors), and the decomposition is done in
 411 two dimensions. In order to avoid issues of buffer overflow and deadlock, we use
 412 `isend()`. The subsequent `waitall` covers the `isend` calls to all neighbors. Note
 413 that many lines of ancillary code may occur between the posting of the `isend`
 414 operations and the call to the balancing `waitall`. Without automatic transfor-
 415 mation capabilities those program section will have to be manually adjoined as
 416 well.

417 To apply our recipe to the `waitall` operation requires the insertion of the
 418 `awaitall`. Because of the aforementioned symmetry of the communication in
 419 practice all processes behave the same and we can illustrate the adjoint com-
 420 munication by the condensed MPI-CFG shown in Fig. 10(left). On the right we
 421 show the wrapper routines inserted into the code in place of the original calls.
 422 To reach a correct solution we again consider inverting the edge direction in
 423 the communication pattern made symmetric by the insertion of the `awaitall`.
 424 Consequently the `recv` is transformed into `isend`, and the `isend` into a `recv`.
 425 Regarding the `recv` turned `isend` we could either impose restrictions on the
 426 `recv` buffer that are identical to the restrictions imposed on an `irecv` buffer
 427 in the `awaitall` - `waitall` section or accept the spatial overhead of using a
 428 temporary buffer instead. Because the restrictions would have required consid-
 429 erable code changes we employed the temporary buffer option. As part of the
 430 symmetric pattern the user code passes a request parameter to the originally
 431 blocking call. The primary reason is of course the need to accommodate the
 432 passing of the actual request in the adjoint but one will observe that adding
 433 these parameters to the interface reflects the very same symmetry that is the
 434 basis of our adjoining recipe. At this point it may be worthwhile to point out
 435 that the superficially similar sequence `irecv` - `send` - `waitall` would *not* permit
 436 bypassing additional restrictions by means of introducing a temporary buffer.
 437 Here, the `irecv` as the adjoint counterpart of the `send` will have to rely on the
 438 restrictions outlined in Sec. 3.4 which in essence in the original code permit a re-
 439 placement of the `send` with an `isend`. Clearly, this “limitation” has to weighed
 440 against the less efficient fall back option of manually splitting the `waitall` call
 441 into individual `waits` on the one hand, or writing the code in \mathcal{P} to satisfy the
 442 `isend` restriction which in turn can improve communication performance in \mathcal{P} .

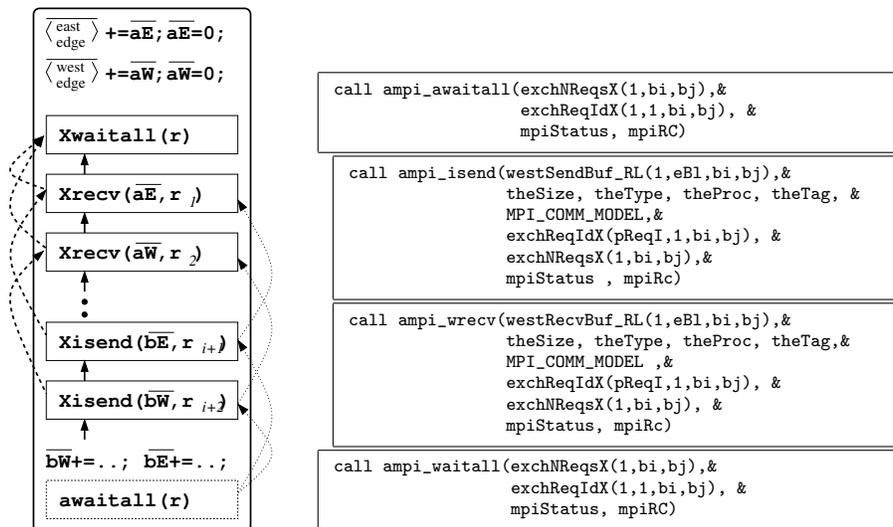


Figure 10: The condensed MPI CFG for the adjoint (left) of the fully symmetric and periodic data exchange uses wrapper routines that encapsulate buffer operations here labeled `Xisend`, `Xrecv`, and `Xwaitall`. The user code is shown in the snippets on the right.

443 The fact that the pairing of `isends` with `irecvs` is not only preferable from the
 444 overall message-passing performance point of view but also permits an easier
 445 program transformation is a rather neat confluence of concepts.

446 The additional buffer parameters are not strictly necessary. The wrapper
 447 could internally associate requests with buffers. On the other hand, the param-
 448 eters are a simple reminder to the user how far the scope of the buffer must
 449 extend. The wrapping approach permits a source transformation with a simple
 450 recipe that directly applies to the wrapped calls and does not require additional
 451 pragma information. Consequently it does not have the same utility for MPI-
 452 aware data-flow analysis. Aside from the extra subroutine call, another source
 453 of overhead is the need to retain separate receive buffers.

454 5 Related Work and Outlook

455 Most of what has been published regarding message passing in the automatic
 456 differentiation context relates to the conceptually simpler forward mode starting
 457 with [12]. The correct association between program variables and their respec-
 458 tive derivatives under MPI might be considered a negligible implementation
 459 issue but has been a practical problem for the application of automatic differen-
 460 tiation in the past [2, 13] and is an issue for the adjoint transformation as well.
 461 Regarding the adjoint mode in particular the description either restricts itself to
 462 plain `send - recv` pairs [3, 4] or describe the hand-written program sections that

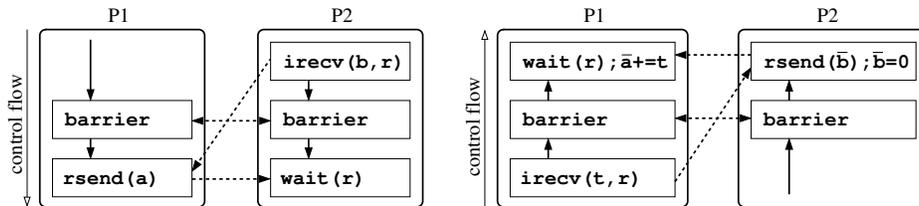


Figure 11: The standard requires that a `recv` has to be posted by the time `rsend` is called, which typically necessitates a `barrier` (left). The adjoint pattern (right) requires a placeholder for inserting the `wait` in P1.

463 “manually” adjoin the communication [9, 10] without an automatic generation
 464 concept for more sophisticated communication patterns.

465 The aim of our paper is to show an approach to the programming of message-
 466 passing logic that guarantees an automatic adjoint transformation can be carried
 467 out by applying rules to subroutine calls. Whether the rules are identified
 468 by pragmas or by a specific set of modified interfaces is an implementation
 469 issue. We have as of yet not covered all communications patterns supported
 470 by the MPI standard. One frequently used MPI call is `barrier`, for instance
 471 in the context of an `rsend`, see Fig. 11. In a logically correct program we
 472 can leave the `barrier` call in place for the adjoint transformation. For the
 473 adjoint the `barrier` call stays in place, however, a vertex transformation recipe
 474 again requires context information and a nonoperational counterpart to the
 475 `rsend` to make the pattern symmetric, e.g. we can introduce an *anti rsend*
 476 or an appropriate communication channel pragma. We cannot claim to have
 477 a prototype with complete coverage of all constructs provided by the current
 478 MPI standard. However, just like the MPI standard itself evolves to meet
 479 user demands we can expand the coverage of an adjointable MPI paired with
 480 automatic differentiation tools based on the techniques explained in this paper.
 481 The prototype implementation done for the MITgcm use case can serve as a
 482 starting point but reaching a consensus among the main tool developers how an
 483 adjointable MPI should be implemented is the eventual goal.

484 6 Summary

485 We provide the rationale for automating the adjoint transformation of message-
 486 passing programs with the need for efficient gradient computation afforded by
 487 automatic differentiation and the difficulty of achieving this goal by other means.
 488 The paper discusses the options for automatically generating an adjoint program
 489 for frequently used communication patterns in message-passing programs. We
 490 show necessary and sufficient requirements to ensure a subroutine call based
 491 set of transformation recipes yields a correct result. The basis for deriving
 492 the recipes are communication graphs. The adjoining semantics requires the
 493 inversion of the communication edge direction and we need to keep the result-

494 ing program deadlock free and efficient. To achieve both goals we introduce
495 additional edges and vertices which make the communication graph symmetric
496 with respect to the edge direction. A PGAS representation of the basic
497 communication patterns provides the basic framework to prove the correctness
498 of the transformations we propose and we show the proving technique on one
499 particular pattern. The automatic transformation tool has to be able to recog-
500 nize the communication calls participating in a particular pattern. Because we
501 want to guarantee the automatic adjoinability of the message-passing program
502 in question we do not want to rely on program analysis that may or may not
503 be able to discern the patterns correctly. Instead, we propose to have the ap-
504 plication programmer either distinguish patterns by means of pragma-identified
505 communication channels or via using a set of specific wrapper routines that dis-
506 tinguish message-passing operations otherwise identical in MPI based on their
507 pattern context. Compared to the alternative of having to hand-code the ad-
508 joint communication the added effort required from the application programmer
509 is rather negligible. Pursuing the approach of identifying communication chan-
510 nels permits improved data flow analysis and opens opportunities of program
511 modifications beyond the generation of adjoints. A use case for our approach
512 was the communication logic implemented in the MITgcm ocean model. We
513 demonstrated the ability to replace the hand-written adjoint communication
514 layer with an automatically generated one. Our future work will concentrate
515 on exploring the implementation options and provide a comprehensive solution
516 that can be used with multiple automatic differentiation tools.

517 Acknowledgments

518 We would like to thank Bill Gropp, Rusty Lusk, Rajeev Thakur and Darius
519 Buntinas for providing insights into the MPI standard rationale and the MPICH
520 implementation. Hovland and Utke are supported by the Mathematical, Informa-
521 tion, and Computational Sciences Division subprogram of the Office of Ad-
522 vanced Scientific Computing Research, Office of Science, U.S. Dept. of Energy
523 under Contract DE-AC02-06CH11357. Hill and Utke are partially supported
524 by the NASA Modeling Analysis and Prediction Program.

525 References

- 526 [1] E. Boman, D. Bozdag, U. Catalyurek, K. Devine, A. Gebremedhin, P. Hov-
527 land, A. Pothen, and M. Strout. Enabling high performance computa-
528 tional science through combinatorial algorithms. In *SciDAC 2007, J. Phys.:*
529 *Conf.Ser.*, volume 78, page 012058(10pp), Bristol, Philadelphia, 2007. IOP.
- 530 [2] Alan Carle and Mike Fagan. Automatically differentiating MPI-1
531 datatypes: The complete story. In Corliss et al. [6], chapter 25, pages
532 215–222.

- 533 [3] C.Faure, P.Dutto, and S.Fidanova. Odysée and parallelism: Extension and
534 validation. In *Proceedings of The 3rd European Conference on Numerical*
535 *Mathematics and Advanced Applications, Jyväskylä, Finland, July 26-30,*
536 *1999*, pages 478–485. World Scientific, 2000.
- 537 [4] B. Cheng. A duality between forward and adjoint MPI communication
538 routines. In *Computational Methods in Science and Technology*, pages 23–
539 24. Polish Academy of Sciences, 2006.
- 540 [5] Cristian Coarfa, Yuri Dotsenko, John Mellor-Crummey, François Canton-
541 net, Tarek El-Ghazawi, Ashrjit Mohanti, Yiyi Yao, and Daniel Chavarría-
542 Miranda. An evaluation of global address space languages: co-array for-
543 tran and unified parallel c. In *PPoPP '05: Proceedings of the tenth ACM*
544 *SIGPLAN symposium on Principles and practice of parallel programming*,
545 pages 36–47, New York, USA, 2005. ACM.
- 546 [6] George Corliss, Christèle Faure, Andreas Griewank, Laurent Hascoët, and
547 Uwe Naumann, editors. *Automatic Differentiation of Algorithms: From*
548 *Simulation to Optimization*, Computer and Information Science, New York,
549 2002. Springer.
- 550 [7] S. Dutkiewicz, M.J. Follows, P. Heimbach, and J. Marshall. Controls on
551 ocean productivity and air-sea carbon flux: An adjoint model sensitivity
552 study. *Geophys. Res. Let.*, 33:L02603, 2006.
- 553 [8] Andreas Griewank. *Evaluating Derivatives: Principles and Techniques of*
554 *Algorithmic Differentiation*. Number 19 in Frontiers in Appl. Math. SIAM,
555 Philadelphia, 2000.
- 556 [9] P. Heimbach, C. Hill, and R. Giering. Automatic generation of efficient
557 adjoint code for a parallel Navier-Stokes solver. In J.J. Dongarra, P.M.A.
558 Sloot and C.J.K. Tan, editor, *Computational Science – ICCS 2002*, volume
559 2331 of *Lecture Notes in Computer Science*, pages 1019–1028. Springer-
560 Verlag, Berlin (Germany), 2002.
- 561 [10] Chris Hill, Alistair Adcroft, Daniel Jamous, and John Marshall. A strat-
562 egy for terascale climate modeling. In *Proceedings of the Eighth ECMWF*
563 *Workshop on the Use of Parallel Processors in Meteorology*, pages 406–425.
564 World Scientific, 1999.
- 565 [11] Chris Hill, Véronique Bugnion, Mick Follows, and John Marshall. Eval-
566 uating carbon sequestration efficiency in an ocean circulation model by
567 adjoint sensitivity analysis. *J. Geophysical Research*, 109(C11005), 2004.
568 doi:10.1029/2002JC001598.
- 569 [12] P. Hovland. *Automatic Differentiation of Parallel Programs*. PhD thesis,
570 University of Illinois at Urbana-Champaign, 1997.

- 571 [13] Paul D. Hovland and Christian H. Bischof. Automatic differentiation of
572 message-passing parallel programs. In *Proceedings of the First Merged In-*
573 *ternational Parallel Processing Symposium and Symposium on Parallel and*
574 *Distributed Processing*, pages 98–104, Los Alamitos, CA, 1998. IEEE Com-
575 puter Society Press.
- 576 [14] J. Marotzke, R. Giering, K. Q. Zhang, D. Stammer, C. Hill, and T. Lee.
577 Construction of the Adjoint MIT Ocean General Circulation Model and Ap-
578 plication to Atlantic Heat Transport Sensitivity. *J. Geophysical Research*,
579 104(C12):29,529–29,547, 1999.
- 580 [15] J. Marshall, C. Hill, L. Perelman, and A. Adcroft. Hydrostatic, quasi-
581 hydrostatic and nonhydrostatic ocean modeling. *J. Geophysical Research*,
582 102, C3:5,733–5,752, 1997.
- 583 [16] MIT general circulation model. <http://mitgcm.org>. source code, docu-
584 mentation, links.
- 585 [17] Uwe Naumann, Laurent Hascoët, Chris Hill, Paul Hovland, Jan Riehme,
586 and Jean Utke. A framework for proving correctness of adjoint message
587 passing programs. Technical Report AIB-2008-06, RWTH Aachen Uni-
588 versity, 2008. accepted for proceedings of EuroPVM/MPI 2008; online at
589 <http://aib.informatik.rwth-aachen.de/2008/2008-06.ps.gz>.
- 590 [18] D. Shires, L. Pollock, and S. Sprenkle. Program flow graph construction
591 for static analysis of MPI programs. In *International Conference on Paral-*
592 *lel and Distributed Processing Techniques and Applications (PDPTA 99)*,
593 1999.
- 594 [19] Marc Snir and Steve Otto. *MPI - The Complete Reference: The MPI Core*.
595 MIT Press, Cambridge, MA, USA, 1998.
- 596 [20] D. Stammer, C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke,
597 A. Adcroft, C. Hill, and J. Marshall. Volume, heat, and freshwater trans-
598 ports of the global ocean circulation 1993–2000, estimated from a gen-
599 eral circulation model constrained by World Ocean Circulation Experiment
600 (WOCE) data. *J. Geophysical Research*, 108(C1):3007–3029, 2003.
- 601 [21] Jean Utke, Uwe Naumann, Mike Fagan, Nathan Tallent, Michelle Strout,
602 Patrick Heimbach, Chris Hill, and Carl Wunsch. OpenAD/F: A modu-
603 lar, open-source tool for automatic differentiation of Fortran codes. *ACM*
604 *Transactions on Mathematical Software*, 34(4), 2008.

605 The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.