Optimizing Remote Communication in X10

ARUN THANGAMANI and V. KRISHNA NANDIVADA, IIT Madras

X10 is a partitioned global address space programming language that supports the notion of *places*; a place consists of some data and some lightweight tasks called activities. Each activity runs at a place and may invoke a place-change operation (using the *at*-construct) to synchronously perform some computation at another place. These place-change operations can be very expensive, as they need to copy all the required data from the current place to the remote place. However, identifying the necessary number of place-change operations and the required data during each place-change operation are non-trivial tasks, especially in the context of irregular applications (like graph applications) that contain complex code with large amounts of cross-referencing objects—not all of those objects may be actually required, at the remote place. In this article, we present AT-Com, a scheme to optimize X10 code with place-change operations.

AT-Com consists of two inter-related new optimizations: (i) AT-Opt, which minimizes the amount of data serialized and communicated during place-change operations, and (ii) AT-Pruning, which identifies/elides redundant place-change operations and does parallel execution of place-change operations. AT-Opt uses a novel abstraction, called *abstract-place-tree*, to capture place-change operations in the program. For each place-change operation, AT-Opt uses a novel inter-procedural analysis to precisely identify the data required at the remote place in terms of the variables in the current scope. AT-Opt then emits the appropriate code to copy the identified data-items to the remote place. AT-Pruning introduces a set of program transformation techniques to emit optimized code such that it avoids the redundant place-change operations. We have implemented AT-Com in the x10v2.6.0 compiler and tested it over the IMSuite benchmark kernels. Compared to the current X10 compiler, the AT-Com optimized code achieved a geometric mean speedup of 18.72× and 17.83× on a four-node (32 cores per node) Intel and two-node (16 cores per node) AMD system, respectively.

CCS Concepts: • **Computing methodologies** → **Parallel programming languages**; **Distributed programming languages**; • **Software and its engineering** → **Compilers**;

Additional Key Words and Phrases: Remote communication, data serialization, program transformation, PGAS languages

ACM Reference format:

Arun Thangamani and V. Krishna Nandivada. 2019. Optimizing Remote Communication in X10. *ACM Trans. Archit. Code Optim.* 16, 4, Article 34 (October 2019), 26 pages.
<https://doi.org/10.1145/3345558>

This journal submission is an extended and reorganized version of a previous article from conference proceedings [28]. Specifically, Sections 1, 5, 6, 7, and 8 include research results that are not published elsewhere (new material: more than 30%). The work has been partially supported by the SERB core research grant number CRG/2018/002488.

Authors' addresses: A. Thangamani, Dept of CSE, IIT Madras, Chennai; email: arunt@cse.iitm.ac.in; V. K. Nandivada, Dept of CSE, IIT Madras, Chennai; email: nvk@iitm.ac.in.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1544-3566/2019/10-ART34

<https://doi.org/10.1145/3345558>

1 INTRODUCTION

With the rapid advancement of many-core systems, it is becoming important to efficiently perform computations in models where the memory may be distributed. X10 [26] is a parallel programming language that uses the partitioned global address space (PGAS) model and provides support for task parallelism. Importantly, X10 supports the distribution of data and computation across shared and distributed memory systems. X10 uses the abstraction of *places*, where each place has some local data (created at that place) and one or more associated activities performing computation over the local data. To access remote data, the activity has to perform a place-change operation (using an *at-construct*). While such expressiveness aids in programmability and data distribution, it may lead to significant communication overheads. We explain the same using a motivating example.

Figure 1(a) shows a code snippet in X10 of the MST (builds a minimum spanning tree) kernel from IMSuite [14]; the `setChildSignal` function checks if any child can start processing in parallel. A child ready to start processing would have already set its corresponding element in the distributed Boolean array `this.setCheck`. The *at* expression checks if a node `i` has set `setCheck(i)`; this value is stored in `atVal`. If any node has set it, then the function returns `true` (code not shown).

X10 supports two main types of non-primitive data: distributed arrays (distributed across one or more places at the time of creation) and non-distributed objects (need to be sent to a remote place, if referred at that place). Consequently, for the code shown in Figure 1(a), the X10 compiler emits code to serialize and send a message containing a deep copy of the complete `this` object and a pointer to the code (`setCheck(i)`). At the remote location, the compiled code will deserialize the message, build a copy of the `this` object, and evaluate the expression. In general, this sending of remote data may incur a significant amount of overhead, and since in Figure 1(a) this remote communication happens inside a loop, the overall cost increases with the number of iterations of the loop.

However, it may be noted that in the code shown, only the `setCheck` field of `this` is getting de-referenced to evaluate the expression; whereas the X10 compiler copies and transmits the complete `this` object, which has many other fields. Note that the current X10 compiler handles distributed arrays efficiently and does not require further optimizations. For example, while copying `this.setCheck` (as part of copying `this`) it only copies the remote-reference of `setCheck`.

Similarly to Figure 1(a), Figure 1(b) shows a snippet of the BF (breadth first search) kernel of IMSuite. The shown snippet starts the BFS creation with the root setting the distance for its neighbors. The function `setDist` (code not shown) uses the distributed array `this.nSet`. Like before, at the *at-construct*, the X10 compiler emits code to transmit the complete `this` object, though only the references to `nSet` is required; note that the sub-partition of `nSet` owned by the target place is already present at the respective target places. We have studied many distributed kernels and found that the amount of such redundantly copied data can be prohibitively large. For example, the X10 compiled MST and BF kernels (snippets in Figure 1, input size=256) copied 76.5GB and 3.0GB data, respectively; we found that a large portion of this data is unused and need not be copied.

In general, it is not trivial to identify the precise data to be copied, especially in the presence of nested *at-constructs*, complex deep procedure calls and arbitrary operations involving heap. As another example, depicting the issue of communication overheads incurred during the translation of code with *at-constructs*, Figure 1(c) shows a code snippet of the HS kernel (to elect a leader in a bidirectional ring network) from IMSuite [14]. The `elect` function starts its computation on the input nodes to select a leader by (i) iterating through all the elements (*points*) in the domain of the `DistributionD` and (ii) performing as many place-change operations. For example, the X10 compiled MST and HS kernels (snippets in Figure 1(a) and (c), input size=256), led to 193,103 and 400,477 remote-communications, respectively. However, the computation can be performed with fewer place-change operations by performing one place-change operation per place and invoking

<pre>def setChildSignal():boolean {<i>//tell children to start</i> for(i in D){ var atVal:boolean =at(D(i))setCheck(i); ... } ... }</pre>	<pre>def Start() { ... at(D(p)){ for(j=0;j<nSet(p).neig. size;j++) { var k:Point = ... setDist(k,nSet(p).d+1); }} }</pre>	<pre>def elect(var ph:Long){ finish { for (i in D) { async at(D(i)) { if(nSet(i).status){...} if(!lValue.equals(0))... }} ... }</pre>
(a) MST: min spanning tree	(b) BF: breadth first search	(c) HS: leader_elect_hs

Fig. 1. Snippets from three IMSuite kernels. D is the distribution of the arrays.

the code body for each point assigned to that place. We have found such similar opportunities in many distributed kernels.

In this article, we present a new scheme (called AT-Com) to efficiently translate X10 code with at-constructs. AT-Com includes two inter-related optimizations: AT-Opt, which identifies/elides the unnecessary data that are otherwise communicated to the remote places, and AT-Pruning, which identifies/elides the redundant place-change operations.

The crux of the proposed AT-Opt optimization is a novel inter-procedural, summary-based, flow-sensitive analysis that precisely tracks the communication across places in terms of the created objects. Unlike other prior works [1, 2], that reason about the places, we do not need to depend on global-value numbering, as every place-change operation has to be handled as an independent place-change operation. Once AT-Opt identifies the required data in terms of the variables in the present scope, it modifies the X10 input program such that only the required data is copied to the remote places. This significantly reduces the remote communication and (consequently) execution time. For example, for MST and BF, AT-Opt reduces the amount of copied data from 76.5 to 10.3GB and 3.0 to 0.13GB, respectively. This in turn, at four places, leads to a speedup of 6.6 $\times$ and 57.3 $\times$, respectively, on a four-node (32 cores/node) Intel system.

Our proposed optimization AT-Pruning reduces redundant place-change operations across places. This leads to a significant reduction in place-change operations and, consequently, execution time. For example, for the MST and HS kernel, AT-Pruning reduces the number of place-change operations from 193,103 to 118,643 and 400,777 to 7,777 respectively. AT-Pruning also emits code to execute place-change operations in parallel, wherever possible. Thus, AT-Pruning reduces the amount of data serialized and the cost of place-change operations. For example, for four places, for the MST and HS kernels, AT-Pruning leads to a speedup of 2.5 $\times$ and 8.5 $\times$, respectively, on a four-node (32 cores/node) Intel system.

Barik et al. [6] present a related scheme to reduce the data communicated during place-change operations by doing scalar replacement. Though their scheme is interesting, its impact is limited in irregular benchmark kernels (like the IMSuite kernels) that have many cross-referencing objects. For example, in Figure 1, the scalar-replacement scheme of Barik et al. cannot be applied. This is because (i) they only focus on scalar fields; setCheck and nSet are distributed arrays, not scalars. (ii) If an at-construct calls a method m by passing an object as an argument or receiver, then the field accesses of that object in m , or in the at-construct after the call to m , cannot be scalar replaced (Figure 1(b)). Note that though setCheck(i) and nSet(p).neig.size are scalars, they cannot be scalar replaced, as the associated arrays are distributed across multiple places and cannot be dereferenced without performing a place-change operation.

Though we discuss AT-Com in the context of X10, it can also be applied to other PGAS languages like HJ [15] and Chapel [8]. Similarly to X10, these languages also support the abstractions/constructs like places/at-constructs and while executing a place-change operation, the reachable non-distributed data are required to be copied to the target place.

Our Contributions:

- We propose a novel analysis to track the flow of objects across places. We are not aware of any other prior work that does so for minimizing communication overheads.
- We propose a new scheme AT-Com to optimize X10 programs that perform remote communication. It includes two optimizations: (i) AT-Opt, which avoids the copying of redundant data across places (Sections 3 and 4), and (ii) AT-Pruning, which removes redundant place-change operations and does parallel execution of place-change operations (Section 5).
- We extend AT-Com to handle programs that may throw exceptions (Section 6).
- We have implemented AT-Com in the x10v2.6.0 compiler C++ backend.
- We have evaluated AT-Com over all the IMSuite kernels on two different hardware systems: a four-node (32 cores/node) Intel system and a two-node (16 cores/node) AMD system. Our evaluations show that compared to the baseline x10v2.6.0 compiler, AT-Com leads to geometric mean speedups of 18.72 $\times$ and 17.83 $\times$ on the Intel system and AMD system, respectively (Section 8).

2 BACKGROUND

In this section, we briefly discuss a few X10 constructs and some pertinent X10 concepts. Interested readers may refer to the X10 specification [26] for details.

- `async S` : spawns a new asynchronous activity to execute `S`.
- `finish S` : acts as join point and waits for all spawned activities in `S` to terminate.
- `at(P) S` : the place-change operator is a synchronous construct and executes the statement `S` at place `p`. The syntax `at(p) async S` spawns an activity at place `p` to execute statement `S`. For the ease of presentation, we represent each such `at`-construct using the sequence of three instructions: `at-entry; S; at-exit`.
- `Distribution` : The language provides the abstraction of a point in n -dimensional space. A set of points can be distributed across the places. Assuming `D` is a distribution, the loop “`for (i in D) S`” iterates over the points in the domain of `D`, ordered by the places.

In X10, the implementation of an `at`-construct of the form “`at(p) S`” involves sending the serialized data needed to execute `S` to the remote place `p` and deserializing the data at `p`. To determine the required data, the compiler analyzes `S` and identifies the referenced variables and sends across all the non-distributed data reachable from them. Each place is assigned with a unique id. The program execution starts at the place 0.

At runtime, the hosts, initial count for places and workers can be set by using the environment variables `X10_HOSTFILE` (or `X10_HOSTLIST`), `X10_NPLACES`, and `X10_NTHREADS`, respectively.

3 AT-OPT: OPTIMIZING REMOTE DATA TRANSFERS

In this section, we propose a compile-time technique AT-Opt to optimize X10 programs that use `at`-constructs. During the compilation of each `at`-construct, the existing X10 compiler emits code to conservatively serialize all the objects (and variables of primitive type) that may be referred to at the remote-place. As discussed in Section 1, this leads to significant overheads. AT-Opt reduces these overheads. For each `at`-construct, AT-Opt conservatively identifies the data “required” at the remote-place (in terms of the local variables, and the reachable fields thereof, in the current scope) and emits code to send/receive only that data. For simplicity, in this section, we assume that the input programs do not throw exceptions. In Section 6, we discuss how we handle exceptions.

AT-Opt has two main phases: (1) Analysis phase (to identify the required data) and (2) Code generation phase (to emit the optimized code). We now discuss both of these phases.

L	= Set of all the program labels.	N_o	= Set of all the abstract-objects.
$L_p \in L$	= Set of labels of the at-constructs.	N_v	= Set of all the variables.
$L_o \in L$	= Set of labels of the new-statements.	N_p	= Set of all the abstract places.

Fig. 2. Definitions of different sets.

$POC : N_o \rightarrow N_p$	place of creation	$\left. \begin{array}{l} MayWS_j \\ MustWS_j \end{array} \right\} \subseteq N_o \times Fields$	$\left\{ \begin{array}{l} \text{written before } L_j, \\ \text{at current place} \end{array} \right.$
$pOf : L \rightarrow N_p$	place of statement		
$RS_j \subseteq N_o \times Fields$	$\left\{ \begin{array}{l} \text{read before } L_j; \text{ but} \\ \text{defined at ancestor place} \end{array} \right.$	$AAL_j \subseteq N_v \cup (N_o \times Fields)$	$\left\{ \begin{array}{l} \text{Ambiguous access} \\ \text{list} \end{array} \right.$

Fig. 3. Auxiliary data structures.

3.1 AT-Opt Analyzer

For each function, in the input program, the AT-Opt analyzer creates two graphs: (1) an *abstract-place-tree* that captures the place-change operations (from a “source” to “target”) and (2) a flow-sensitive *points-to graph* that captures the points-to information of X10 objects (by extending the escape-to connection-graph described by Agarwal et al. [1]). We first elaborate on these two graphs.

3.1.1 Abstract-place-tree (APT). For each function in the input X10 program, an APT defines the relationship among the instances of different at-constructs in the function. Each at-construct corresponds to one or more place-change operations at runtime. Say the set of labels¹ of these constructs is given by L_p (see Figure 2). An APT is defined by the pair (N_p, E_p) , where $N_p = \{p_i | L_i \in L_p\}$. Thus, each $p_i \in N_p$ represents an abstract place-change operation. Given two nodes $p_i, p_j \in N_p$, we say that $(p_i, p_j) \in E_p$ if L_j is present in the body of p_i . An interesting aspect of the APT data structure is that it exposes the dataflow between places, as per the X10 semantics: Changes done to any data structure (not a global reference or a distributed array) at a place node are not visible to its ancestors and siblings. That is, in the APT, data flow only from the parent to the children.

3.1.2 Points-to Graph (PG). We use the definitions in Figure 2 and a special object node $O_\top$ (that represents all the non-analyzable abstract-objects in the program) to define a points-to graph. A points-to graph is a directed graph (N, E) , where $N = N_o \cup N_v \cup \{O_\top\}$. Similarly to the discussion of APT, each abstract object $o \in N_o (= \{o_i | L_i \in L_o\})$ may represent one or more instances of objects created at the corresponding labels at runtime. We call an abstract object that represents multiple runtime object instances a *summary object*.

The set E comprises two types of edges:

1. *points-to edges* $\subseteq N_v \times N_o \cup \{O_\top\}$: These edges of the form $v \rightarrow^p o$ are created because of assignment of objects (say, o) to variables (say, v).
2. *field edges* $\subseteq N_o \cup \{O_\top\} \times Fields \times N_o \cup \{O_\top\}$: These edges of the form $o_1 \rightarrow^{fg} o_2$ are created because of assignment of objects (say, o_2) to the fields (say, g) of objects (say, o_1).

We call an edge $v \rightarrow^p o$ (or $o_1 \rightarrow^{fg} o_2$) to be a *weak-edge*, if $\exists o' \neq o$, such that $v \rightarrow^p o' \in E$ (or $o_1 \rightarrow^{fg} o' \in E$). Otherwise, we call it a *strong-edge*. We use this classification later in this section to mark objects that can be tracked precisely.

Besides maintaining APT (global) and PG (at each statement), we maintain a few other data structures, listed in Figure 3. POC returns the place of object-creation and pOf returns the

¹Without any loss of generality, we assume that the input is a simplified X10 program in three-address-code form [19], each statement has a unique associated label, and variables have unique names.

place-node of each statement. For each function g , we assume that all the statements not contained inside any at-construct are executed at the special place p_g . A pair $\langle o_i, f \rangle \in RS_j$ indicates that the field f of o_i is used (read) at a predecessor of L_j at place $\text{pOf}(j)$, and the definition reaching this use is present in one of the APT-ancestors of $\text{pOf}(j)$. At each label L_j , we maintain two “write-sets”: $MayWS_j$ and $MustWS_j$ to hold the may and must information indicating that the object-field may-/must-be defined at a predecessor of L_j , at $\text{pOf}(j)$. Note: $MustWS_j \subseteq MayWS_j$. An entry $k \in AAL_j$ indicates that k (a variable or an obj-field pair) has some weak-edges in the PG_j and k is accessed at $\text{pOf}(j)$; this set is used for identifying ambiguous objects during code-generation phase (Section 3.2). Note that we separately compute the AAL set instead of deriving it from the points-to graph after the first pass, as otherwise it will lead to imprecision. For example, consider a points-to graph like the one shown in Figure 7(b) and say the object o_j is accessed at the target place. Now, if the points-to graph is used to derive the AAL, then it would include both a , and b , even though the access might be via only the variable b (a strong edge).

3.1.3 Intra-procedural Flow-sensitive Analysis. We now discuss our scheme to perform a flow-sensitive iterative dataflow analysis to build the APTs (global), the points-to graphs (at each label) and the auxiliary data structures, in a combined manner. The points-to-graph construction is standard and is shown for completeness. For the ease of presentation, we now focus just on the intra-procedural component of the analysis. We discuss the inter-procedural extension in Section 4.

Initialization. For each function bar , at the first instruction: (1) APT is initialized to a single root node ($p_{bar} \in N_p$); (2) PG is initialized to (N_i, E_i) , where $N_i = N_v \cup \{O_\top\}$. In bar , for each function parameter $a_j \in N_v$ (including the 0^{th} argument *this*), conservatively, we include an edge $a_j \rightarrow^p O_\top$ in E_i . Also, we add an edge $O_\top \rightarrow^{f,*} O_\top$ to indicate that all field edges from $O_\top$ will point to $O_\top$. The rest of the auxiliary data structures are initialized to empty.

Statements and related operations. Figure 4 shows how we update the APT, PG, and auxiliary data structures on processing the different X10 statements. For each statement $L : Stmt$, each transformation is shown as a transition of the form $X \Rightarrow X'$, where X is a data structure before processing the statement $Stmt$ at label L and X' is the updated data structure after the processing. Unless otherwise specified, each data structure is copied (cloned) to the next statement.

The statements that are of interest to our analysis are as follows:

(i) L : at-entry(p) (ii) L : $a = \text{new } T()$; (iii) L : $a = b$; (iv) L : $a = b.f$; (v) L : $a.f = b$; (vi) L : $a = x.bar(b)$; and (vii) L : at-exit. We now discuss how the processing of each of these statements updates APT, PG, and the other auxiliary data structures.

Entering at L_j : at-entry(p): We create a new place node p_j and add an edge from the current place (given by $\text{pOf}(L_j)$) to the target place (p_j) in APT (Figure 4(a)). Further, we reset $AAL = MayWS = MustWS = RS = \Phi$ (not shown). Note: Only the at-entry instruction updates the APT.

Exiting at L_j : at-exit(L_i): We restore PG, RS, $MayWS$, $MustWS$, and AAL to their values at L_i , which has the corresponding at-entry instruction.

Allocation statement. $L_j : a = \text{new } T()$: Besides updating the PG (Rule 1, Figure 4(b)), creates a new object node o_i and updates the points-to edges, we add $(o_i, \text{pOf}(L_j))$ to POC.

Copy statement. $L_j : a = b$: Besides updating the PG (Rule 2, Figure 4(b)): If b has weak-edges, then we update AAL to include b , since b is used here.

Load statement. $L_j : a = b.g$: Besides updating the PG (Rule 3, Figure 4(b)), we update RS, keeping in mind that no definitions from the current place are added to RS. If b or $b.g$ has weak-edges, then we add them to AAL set appropriately.

Store statement. $L_j : a.g = b$: If a has weak-edges or a points-to a summary node, then we perform a weak update; else we perform a strong update (Rule 4, Figure 4(b)). Besides updating the PG, we add all of the object-field pairs that may be referred to by $a.g$ to $MayWS$. These pairs are also added

$$L_j : \text{at-entry}(p) \ (N_p, E_p) \Rightarrow (N_p \cup \{p_j\}, E_p \cup \{(p\text{Of}(L_j) \rightarrow p_j)\})$$

(a) Impact on the APT

1. $L_j : a = \text{new } T()$	2. $L_j : a = b$
$(N, E) \Rightarrow (N \cup \{o_j\}, (E - \{a \rightarrow^p y \mid a \rightarrow^p y \in E\}) \cup \{a \rightarrow^p o_j\})$	$(N, E) \Rightarrow (N, (E - \{a \rightarrow^p y \mid a \rightarrow^p y \in E\}) \cup \{a \rightarrow^p z \mid b \rightarrow^p z \in E\})$
$POC \Rightarrow POC \cup \{(o_j, p\text{Of}(L_j))\}$	$AAL \Rightarrow AAL \cup \{b\} \ // \ \text{if } b \text{ has weak-edges.}$
3. $L_j : a = b.g$	5. $L_j : a = x.\text{bar}(b)$
$(N, E) \Rightarrow (N, (E - \{a \rightarrow^p y \mid a \rightarrow^p y \in E\}) \cup \{a \rightarrow^p z \mid b \rightarrow^p x \in E \wedge x \rightarrow^{\text{fg}} z \in E\})$	$(N, E) \Rightarrow (N, E \cup \{a \rightarrow^p O_\top\} \cup \{y \rightarrow^{\text{fg}*} O_\top \mid x \rightarrow^+ y \in E\} \cup \{z \rightarrow^{\text{fg}*} O_\top \mid b \rightarrow^+ z \in E\})$
$RS \Rightarrow RS \cup \{\langle o_i, g \rangle \mid b \rightarrow^p o_i \in E \wedge \langle o_i, g \rangle \notin \text{MustWS}\}$	$RS \Rightarrow RS \cup \{\langle o_i, * \rangle \mid x \rightarrow^+ o_i \in E \vee b \rightarrow^+ o_i \in E\}$
$AAL \Rightarrow AAL \cup \{b\} \ // \ \text{if } b \text{ has weak-edges}$	$AAL \Rightarrow AAL \cup \{z \mid z \in \{x, b\} \wedge z \text{ has weak-edges in } E\}$
$AAL \Rightarrow AAL \cup \{\langle o_i, g \rangle \mid b \rightarrow^p o_i \in E \wedge b.g \text{ has weak-edges}\}$	$AAL \Rightarrow AAL \cup \{\langle o_i, g \rangle \mid (x \rightarrow^+ o_i \in E \vee b \rightarrow^+ o_i \in E) \wedge \langle o_i, g \rangle \text{ is a weak-edge}\}$
4(i). $L_j : a.g = b$ (Strong update)	4(ii). $L_j : a.g = b$ (Weak update)
$(N, E) \Rightarrow (N, (E - \{y \rightarrow^{\text{fg}} z \mid a \rightarrow^p y \in E \wedge y \rightarrow^{\text{fg}} z \in E\}) \cup \{y \rightarrow^{\text{fg}} x \mid b \rightarrow^p x \in E \wedge a \rightarrow^p y \in E\})$	$(N, E) \Rightarrow (N, E \cup \{y \rightarrow^{\text{fg}} x \mid b \rightarrow^p x \in E \wedge a \rightarrow^p y \in E\})$
$\text{MayWS} \Rightarrow \text{MayWS} \cup \{\langle o_i, g \rangle \mid a \rightarrow^p o_i \in E\}$	$\text{MayWS} \Rightarrow \text{MayWS} \cup \{\langle o_i, g \rangle \mid a \rightarrow^p o_i \in E\}$
$\text{MustWS} \Rightarrow \text{MustWS} \cup \{\langle o_i, g \rangle \mid a \rightarrow^p o_i \in E\}$	$AAL \Rightarrow AAL \cup \{a\} \ // \ a \text{ might not be in } AAL, \text{ yet}$
$AAL \Rightarrow AAL \cup \{b\} \ // \ \text{if } b \text{ has weak-edges.}$	$AAL \Rightarrow AAL \cup \{b\} \ // \ \text{if } b \text{ has weak-edges.}$

(b) Impact on the points-to-graph PG and auxiliary data structures (only the updated ones are shown). For brevity, we omit the subscripts on the auxiliary data-structures *RS*, *MustWS*, *MayWS*, and *AAL*.

Fig. 4. Rules for intra-procedural analysis.

to *MustWS* if we are performing a strong update. If *b* has weak-edges, then *b* is added to *AAL*. In case of weak-update, we also add *a* to *AAL*, as *a* might not have been added so far to the current *AAL*.

Function call (intra-procedural analysis). $L_j : a = x.\text{bar}(b)$: Here we make conservative assumptions on the impact of the function call on the arguments, receiver and the return value. We first update the PG (Rule 5, Figure 4(b)). We use the notation $p \rightarrow^+ q$ to indicate that *q* is reachable from *p* (in the current PG) after traversing one or more edges (points-to or field). We add all the weak-edges reachable from *x* and *b* to *AAL*. We conservatively assume that all the fields reachable from *x* and *b* are read in the method *bar* and add them to *RS*.

Merge Operation. At each control-flow join point, we merge the *APT*, *PG*, and the auxiliary sets. The merging of graphs is done by taking a union of the nodes and edges. For *RS*, *MayWS*, and *AAL*, we merge the sets by performing the set-union operation. We merge the *MustWS* sets by performing the set-intersection operation.

Termination. We follow the standard iterative dataflow analysis approach [19] and stop after reaching a fixed point (from the point of view of *APT* and *PG*).

Post analysis. Finally, we populate two sets for each node in the *APT*: (i) cumulative read-set *CRS* and (ii) cumulative ambiguous-access-list *CAAL*. A pair $\langle o_i, f \rangle \in \text{CRS}_n$ indicates that the field *f* of the object *o_i* is defined at one of the *APT* predecessors of *n* and that definition may reach a statement in *n* or one of its *APT* successors. An entry $k \in \text{CAAL}_n$ indicates that *k* is in the ambiguous-access-list of either *n* or one of its *APT* successors.

Computation of *CAAL* and *CRS*: We use X_i to refer to the data structure *X* before processing the statement labeled L_i . Note that each node in *APT* can be represented by a pair $\langle i, j \rangle$, where, L_i and L_j are the labels of the first and the last instruction, respectively, of the node. We traverse the *APT* in post-order and for any *APT* node $n = \langle x, y \rangle$ set: (1) $\text{CRS}_n = \text{RS}_y \cup_{\langle p, q \rangle \in \text{aptChild}(n)}$

Fig. 5. (a) Example synthetic X10 code. (b) Generated APT and points-to graph for code in Figure 5(a).

RS_{16}	$\{\langle o_7, r_1 \rangle\}$	$CRS_{\langle 1, 23 \rangle}$	$\{\langle o_2, r_1 \rangle\}$	$MayWS_4$	$\{\langle o_2, r_1 \rangle\}$	$MustWS_4$	$\{\langle o_2, r_1 \rangle\}$
RS_{21}	$\{\langle o_2, r_1 \rangle, \langle o_2, r_2 \rangle\}$	$CRS_{\langle 4, 22 \rangle}$	$\{\langle o_2, r_1 \rangle, \langle o_2, r_2 \rangle, \langle o_7, r_1 \rangle\}$	$MayWS_{11}$	$\{\langle o_2, r_2 \rangle, \langle o_7, r_1 \rangle, \langle o_7, r_3 \rangle\}$	$MustWS_{11}$	$\{\langle o_2, r_2 \rangle, \langle o_7, r_1 \rangle, \langle o_7, r_3 \rangle\}$
POC	$\{\langle o_2, p_f \rangle, \langle o_7, p_4 \rangle\}$			$MayWS_{18}$	$\{\langle o_7, r_3 \rangle\}$	$MustWS_{18}$	$\{\langle o_2, r_2 \rangle, \langle o_7, r_1 \rangle\}$

(a) (b) (c) (d)

Fig. 6. Auxiliary data structures at different program points and nodes of APT.

$(RS_q \cup (CRS_{\langle p, q \rangle} - MustWS_q))$ and (2) $CAAL_n = AAL_y \cup_{\langle p, q \rangle \in aptChild(n)} CAAL_{\langle p, q \rangle}$, where $aptChild(n)$ returns all nodes reachable from n in the APT.

Example. Consider the example code shown in the Figure 5(a). Here, $L_p = \{p_4, p_{11}, p_{18}\}$ and $L_o = \{o_2, o_7\}$. Figure 5(b) shows the generated APT and PGs. For brevity, we only show the contents of the PG just before the at-constructs. Note that the nodes in the APT can also be represented as a pair of indices. For example, p_f can be represented as $\langle 1, 23 \rangle$ and p_4 as $\langle 4, 22 \rangle$. Figure 6 shows the contents of RS , $MayWS$, and $MustWS$ sets at different representative program points. CRS and POC maps are shown as seen after the analysis of the function f . For this example, $AAL = \phi$.

3.2 Optimized Code Generation

We now discuss our code generation scheme that uses the information (APT , PG , RS , CRS , $MayWS$, POC , and $CAAL$) computed in Section 3.1.3 to generate code that copies only the required data to the target places during a place-change operation. For the ease of explanation, we show the rules as a source-to-source translation scheme; the generated code is compiled by the current X10 compiler.

While compiling an at-construct, the current X10 compiler captures all the free variables (including this) that are referenced in the body of the at-construct and emits code to copy all the data reachable from these free variables. For example, in Figure 5(a) for the at-construct at line 11, the whole of object b will be copied to the target place. We take advantage of this approach of the X10 compiler and use a simple scheme to emit optimized code. We first illustrate our scheme using the code in Figure 5(a). We emit code to copy $b.r1$ to a temporary (say $t3$) just before line 11. In the body of the at-construct, we create an empty object (say $b1$, of type B), set $b1.r1 = t3$, and replace every occurrence of $b.r1$ in the body of the at-construct with $b1.r1$. Note: In this generated code, b is not one of the captured free variables, and hence the object pointed-to by b will not be copied by the X10 compiler; $t3$ will be copied instead. We first present a brief discussion on the impact of objects pointed-to by weak-edges and then detail our code generation scheme.

Fig. 7. An example to illustrate “ambiguous” objects.

Ambiguous objects: In a points-to graph, we call an object o_i to be ambiguous if o_i is reachable from some node, say, $x \in PG.N$, and the path from x to o_i in PG contains a weak-edge. Unlike a non-ambiguous object, we cannot precisely determine which of the fields of the ambiguous objects are to be copied. For example, Figure 7 shows a small snippet of X10 code and its points-to graph (before line 2). At runtime line 3 may write to field $\langle o_i, r1 \rangle$ or $\langle o_j, r1 \rangle$, and, hence, the field dereference at line 4 may read the value of the field $\langle o_j, r1 \rangle$ written at line 3 or in place p_1 . Consequently, at compile time we would not know if $b.r1$ should be copied from p_1 . Because of such inexactness during compilation time (and to be sound), we deal with the ambiguous objects (for example, o_i and o_j , in Figure 7) conservatively and copy the full objects (provided, the objects are being dereferenced).

We now describe how our code-generation pass handles the statements discussed in Section 3.1.3. Of these statements, handling the `at`-construct (described first) is more involved than the rest.

Handling `at`-construct L_j :`at-entry`(p). Say the corresponding APT node is $\langle j, k \rangle$. We first compute the set of all the ambiguous objects that are reachable from the elements of $CAAL_{\langle j, k \rangle}$: $AOS = \{o_i | n \in CAAL_{\langle j, k \rangle} \wedge n \rightarrow^{+w} o_i \in PG_j\}$; here $n \rightarrow^{+w} o_i \Rightarrow o_i$ is reachable from n (in PG_j) after traversing one or more weak edges. During a place-change operation, the objects in AOS will be copied fully. This code generation phase for the remaining objects has two parts.

(A) Code emitted immediately before label L_j : Figure 8(a) emits code to save the field $\langle o_i, g \rangle$ that is used in the successor(s) of $pOf(j)$ (in the APT) if (1) $CRS_{\langle j, k \rangle}$ contains $\langle o_i, g \rangle$ and o_i is non-ambiguous; (2) o_i created at $pOf(j)$ or $MayWS_j$ contains $\langle o_i, g \rangle$ (that is, $\langle o_i, g \rangle$ may be written to at the parent place $pOf(j)$); and (3) either $\langle o_i, g \rangle$ is a scalar field or it points to an ambiguous object (Lines 3–7).

As we will see in Figure 8(b), we may create new substitute objects inside an `at`-construct. And in the body of the `at`-construct, whenever there is a reference to the original object, those references have to be replaced by the substitute objects. To aid in this process, we maintain a map H , which takes as input an object o_i and returns the name of the variable pointing to the substitute object. We save the old value of H before processing an `at`-construct (`at-entry` statement, that is) and restore H to the saved value after completing the processing of `at-exit`.

To emit the correct code, we need to identify a variable x that points to the substitute object of o_i (if present), or o_i in PG_j (Lines 8–11). Then we create a temporary (name in $T1$) and emit code to copy $x.g$ to the temporary. We remember this mapping of $\langle o_i, g \rangle$ to $T1$ in a global map $tMap$.

(B) Code emitted in the body of the `at`-construct at L_j : Figure 8(b) emits code to create an object for o_i and re-store the value of $\langle o_i, g \rangle$ from temporaries (added in Step A); we further improve it in Section 7. For each pair $\langle o_i, g \rangle$ read in the body of `at`-construct and o_i is not an ambiguous object, we call the function `createObject`(o_i, H, S) to emit code to create a substitute object for o_i (Line 5); say stored in a variable named tx_0 . We then check if $\langle o_i, g \rangle$ is a non-scalar field pointing to a non-ambiguous object say o_j . If so, then we emit code (Line 7) to create a substitute object for o_j (say, stored in a variable named tx_1) and initialize $tx_0.g$ to tx_1 (Line 8). Otherwise (either $\langle o_i, g \rangle$ is a scalar, or o_j is ambiguous), we lookup (in $tMap$) the temporary (say, named tx_3) in which the value

```

1 Input :  $PG_j, AOS_j, CRS_{(j,k)}, MayWS_j, POC$ .
2 begin
3   foreach  $\langle o_i, g \rangle \in CRS_{(j,k)} \wedge o_i \notin AOS_j$  do
4     if  $POC(o_i) == pOf(j) \vee \langle o_i, g \rangle \in MayWS_j$ 
5       then
6         bool  $f1 = \exists (o_i \rightarrow^{fg} o_l) \in PG_j.E$ ;
7         bool  $f2 = f1 \wedge (o_l \in AOS_j)$ ;
8         if  $\neg f1 \vee f2$  then // Emit code to copy
9           if  $H.contains(o_i)$  then //  $o_i$  is
10            referred to by a new name
11            |  $x = H.get(o_i)$ ;
12          else
13            Set  $x$  to one of the variables
14            pointing-to  $o_i$  in  $PG_j.E$ ;
15            // Three-address-code input.
16            Hence, each object is pointed
17            to by at-least one variable.
18          String  $T1 = \text{new TempName}()$ ;
19          Emit (" $val$  "  $\| T1 \| "=" \| x \| ".g;$ ");
20           $tMap.put(\langle o_i, g \rangle, T1)$  // save
21          mapping

```

(a) Emit the required code before the at-construct of the APT node $\langle j, k \rangle$.

```

1 Input :  $PG_j, AOS, RS_k, MayWS_k$ 
2 begin
3   Set  $S = \phi$ ;
4   foreach  $\langle o_i, g \rangle \in RS_k \wedge o_i \notin AOS$  do
5     createObject( $o_i, H, S$ );
6     if  $\exists (o_i \rightarrow^{fg} o_j) \in PG_j.E \wedge o_j \notin AOS$ 
7       then
8         createObject( $o_j, H, S$ );
9         Emit( $H.get(o_i) \| "." \| g \| "=" \|$ 
10           $H.get(o_j) \| ";"$ );
11       else Emit( $H.get(o_i) \| "." \| g \| "=" \|$ 
12           $tMap.get(\langle o_i, g \rangle) \| ";"$ );
13   foreach  $\langle o_i, g \rangle \in MayWS_k \wedge o_i \notin AOS$  do
14     createObject( $o_i, H, S$ );

```

(b) Emit code to copy data from the temporaries.

```

12 Function createObject ( $o, H, S$ )
13 begin
14   if  $\neg S.contains(o)$  then
15      $S = S \cup o$ ; String  $T1 = \text{new TempName}()$ ;
16     Emit (" $val$  "  $\| T1 \| "=new " \| typeOf(o) \| "("$ ;";");
17      $H.put(o, T1)$ ;

```

(c) Emits code to create a "substitute" obj for o_i .

Fig. 8. Auxiliary functions.

```

1 def f():void{
2   val a:A=new A();
3   a.r1 = ...
4   t1 = a.r1;
5   at (P) {
6     val a1:A=new A();
7     a1.r2 = ...
8     ... = a1.r2;
9     val b:B=new B();
10    b.r1 = ...
11    for(i in D){
12      b.r3 = a1;
13      t3 = b.r1;
14      at(D(i)){
15        val b1:B=new B();
16        b1.r1 = t3;
17        val a2:A=new A();
18        ... = b1.r1;
19        val c:A = a2;
20        c.r1 = ...
21        ... = a2.r1; }}
22    t2 = a1.r2;
23    at(Q){
24      val a3:A=new A();
25      a3.r1 = t1;
26      a3.r2 = t2;
27      ... = a3.r1;
28      ... = a3.r2; }}

```

Fig. 9. Optimized code for Figure 5(a).

of $\langle o_i, g \rangle$ was stored before the at-construct and initialize $tx_0.g$ to tx_3 (Line 9). For each pair $\langle o_i, g \rangle$ written in the body of at-construct and o_i is not ambiguous, call the function $\text{createObject}(o_i, H, S)$ to check and emit code to create a substitute object for o_i (Lines 10 and 11).

Handling statements other than the at-construct. For statements $a=b$, $a=b.f$, $a.f=b$, and $a=x.f(b)$, we check if the objects pointed to by the variables and fields have substitute objects created in the current scope (and are present in the H map), and if so we replace the variables with the temporaries created. Section 7 discusses a further optimization for the statement $a=b$.

Code generation for our running example: For the example shown in the Figure 5(a), our code generation pass takes the computed data structures (shown in Figures 5(b) and 6) and generates code as shown in Figure 9.

As it can be seen, unlike the default X10 compiler that copies the complete object (for example, in Figure 5(a) the objects pointed-to by a , a and b , and a for the at-constructs at Lines 4, 11, and 18, respectively) to the destination place, our proposed AT-Opt copies only the required data (for example, in Figure 9 see lines 4, 13, and 22), creates the required substitute objects therein (for example, in Figure 9 see lines 6, 15, and 24) and initializes substitute objects with the copied data.

Note that at Line 6, we only create a substitute object but do not (need to) emit additional code to initialize any of its fields. In contrast, the substitute objects created at Lines 15 and 24 have some of their fields initialized explicitly, because those fields are explicitly live from remote place and referenced later in the code (see Section 7 for a further optimization in the generated code).

In contrast to AT-Opt, the scalar-replacement technique of Barik et al. [6] will serialize the complete object pointed-to by *a*, as (i) it is copied at Lines 10 and 13 and (ii) one of its fields is written to (*a.r2* at Line 5). Their scheme can scalar-replace *b.r1* if it is a scalar field. Otherwise, their scheme cannot scalar-replace *b.r1* and will serialize the complete object pointed-to by *b*.

4 INTER-PROCEDURAL AT-OPT

In this section, we present the inter-procedural extension to the intra-procedural analysis discussed in Section 3. This is based on an extension to the standard summary-based flow-sensitive analysis [19]. In addition to maintaining the standard summaries for points-to information and iterating to a fixed point, we maintain summaries for *CRS* and *CAAL*. For each function node in the call-graph, we maintain (1) Input summary: gives the summary of the points-to information of the function parameter(s) including the *this* pointer. (2) Output Summary: (A) points-to details of function parameters (as seen at the end of the function) and return value (B) cumulative-read-set: *CRS* as computed for the abstract place corresponding to the function call. (C) Cumulative ambiguous accesses list: *CAAL* as computed for the abstract place corresponding to the function call. As expected, the input and output summaries are conservative and sound.

Our inter-procedural analysis follows a standard top-down approach with additional steps to compute or maintain the specialized summaries under consideration. We now discuss some of the salient points therein.

Representation of Objects: In addition to the label in which the object is allocated, we maintain a (finite) list of labels giving a conservative estimation about the context (call-chain) in which the object is created; this list is referred as the context-list of the object.

Initialization. Unlike the intra-procedural analysis, where the analysis of each function starts with a conservative assumption of its arguments, here the analysis begins with an initial points-to graph representing the summary points-to graph of the arguments.

End of analysis of a function. Once we terminate the analysis of a function *bar*, besides creating a summary for the points-to graph, we set the *CRS* (and *CAAL*) in output summary as the *CRS* (and *CAAL*) of the APT node corresponding to the function *bar*; recall that corresponding to each function, we create a special place node and that is the root of the APT for that function.

Processing the statements. All the statements except that of the function call are handled mostly similar to the way they were handled during the proposed intra-procedural analysis (Figure 4). The only difference is that we use the extended representation for the objects. The newly created object is represented as $o_{\langle i, [0] \rangle}$, where $[0]$ represents the context. If this allocation site is present in a function f_1 , then another function f_2 calls f_1 , and this object is made accessible in f_2 (for example, via a return statement in f_1), then the context-list of the object is updated to reflect the call of f_1 in f_2 .

We now discuss how we process the function-call statement (shown in Figure 10). The main difference in processing is related to the handling of input-summary (IS) and taking into consideration the impact of the output-summary (OS) on the arguments and return value. This process is followed for each of the functions that *bar* may resolve to statically.

Impact on IS and OS of the function *bar*. (i) IS: In points-to graph present in the IS of the function *bar*, the formal parameter corresponding to the actual argument *b* is updated to additionally point to the nodes pointed to by *b*. (ii) OS: Say the APT node for the function *bar* is represented by $\langle m, n \rangle$. (1) For each object $o_{\langle i, [C] \rangle}$ in the merged *PG*, if it is created in the function *bar*, then we append the

$$\begin{aligned}
L_j : a = x.bar(b) \quad (N, E) &\Rightarrow (N \cup \{o_{\langle i, [j, C] \rangle} | o_{\langle i, [C] \rangle} \in OS.PG.N \wedge [C] = [ll, C'] \wedge ll \text{ is a label in } bar\}, \\
&E \cup \{a \rightarrow^P o_{\langle i, [C] \rangle} | f_{ret} \rightarrow^P o_{\langle i, [C] \rangle} \in OS.PG.E\} \\
&\cup \{y \rightarrow^{f,g} o_{\langle i, [C] \rangle} | this \rightarrow^+ y \in IS.PG.E \wedge y \rightarrow^{f,g} o_{\langle i, [C] \rangle} \in OS.PG.E\} \\
&\cup \{z \rightarrow^{f,g} o_{\langle i, [C] \rangle} | farg_b \rightarrow^+ z \in IS.PG.E \wedge z \rightarrow^{f,g} o_{\langle i, [C] \rangle} \in OS.PG.E\} \\
RS &\Rightarrow RS \cup (CRS_{\langle m, n \rangle} - LocalObjects_{bar}) // \text{if the } \langle m, n \rangle \text{ is the APT node for } bar. \\
AAL &\Rightarrow AAL \cup \{z | z \in \{x, b\} \wedge z \text{ has weak-edges in } E\} \\
AAL &\Rightarrow AAL \cup \{\langle o_i, g \rangle | (x \rightarrow^+ o_i \in E \vee b \rightarrow^+ o_i \in E) \wedge \langle o_i, g \rangle \text{ is a weak-edge}\} \\
AAL &\Rightarrow AAL \cup (CAAL_{\langle m, n \rangle} - LocalVars_{bar}) // \text{if } \langle m, n \rangle \text{ is the APT node for } bar.
\end{aligned}$$

Fig. 10. Rules to translate a function-call instruction for inter-procedural analysis.

Fig. 11. Block diagram of AT-Pruning (rules in Figure 12).

label L_j to the context-list C . (2) We set b to point to whatever the corresponding formal parameter of bar is pointing to in the OS. (3) We set a to point to whatever the return value is pointing to in the OS. (4) We merge the entries of $CAAL_{\langle m, n \rangle}$ with the current AAL , after removing the local variables of bar . (5) We update CRS by taking a union of current RS with $CRS_{\langle m, n \rangle}$, after removing the local object pairs of bar .

5 AT-PRUNING: REDUCING THE OVERHEADS OF PLACE CHANGE OPERATIONS

In this section, we propose new program transformation techniques to avoid redundant place-change operations and allow parallel execution of the remaining place-change operations—both of these together lead to reduction in the overheads resulting from the place-change operations. Though AT-Pruning is not directly related to communication optimization, these rules reduce the number of place-change operations and thereby reduce the number of remote-communications. For the ease of presentation, in this section, we assume that the input code does not throw any exceptions. In Section 6, we show how we handle code that may throw exceptions.

Figure 11 shows the overall block-diagram of our proposed optimization phase AT-Pruning. This phase takes as input the Abstract Syntax Tree (AST) of the input X10 program and generates an optimized X10 AST. The optimization phase continuously invokes one of the seven mini-transformations (listed in Figure 12), until no further change is possible. Note that except for the Loop+AT-Fusion rule, the rest can be applied in any order without impacting the final generated code. The rule Loop+AT-Fusion is invoked when no more program transformation is possible from other rules listed in Figure 12. Applying the Loop+AT-Fusion rule on top of the generated transformed code helps improve the code generated from AT-Expr parallelization (rule [B(i)]). For efficiency, we apply these transformations in a bottom-up fashion in the AST. We specify the rules in the form $X \Rightarrow Y$ to indicate that code of the form X gets translated to Y . With each rule, we optionally mention (as conditions) the condition under which the rules can be invoked.

Removing redundant place-change operations. In Figure 12, rule [A] removes redundant place-change operations by generating code that invokes a single place-change operation per place (in contrast to one place-change operation per each point in the distribution D). The inner loop (over the place p) in the generated code iterates over the points assigned to the place p . Let P be

[A] Redundant place-change operation elision	
(i) Serial place-change operations for (i in D) { at(D(i)) { S1; } }	$\Rightarrow$ { for (p in place.Places()) { at(p) { for (i in D(p)) { S1; } } }
(ii) Parallel place-change operations for (i in D) { async at(D(i)) { S1; } }	$\Rightarrow$ { for (p in place.Places()) { at(p) { for (i in D(p)) { async S1; } } }
[B] Parallelization of place-change operations	
(i) AT-Expr parallelization <i>//E1 is independent of S1, S2 and other instances of E1</i> for (i in D) { S1; ... = at(D(i)) E1; <i>//S1 has no at-expression</i> S2; }	$\Rightarrow$ { Val r1 = new Rail[T](D.size()); finish for (i in D) { async r1(i) = at(D(i)) E1; } for (i in D) { S1; ... = r1(i); S2; } }
(ii) AT-Stmt parallelization at(p) { for (i in D(p)) { async S1; } }	$\Rightarrow$ { async { at(p) { for (i in D(p)) { async S1; } } }
[C] Fusion	
(i) Loop-fusion <i>//S1 and S2 are independent of each other</i> for (p in place.Places()) { at(p) S1; } for (p in place.Places()) { at(p) S2; }	$\Rightarrow$ { for (p in place.Places()) { at(p) S1; at(p) S2; } }
(ii) AT-fusion at(p) S1; at(p) S2;	$\Rightarrow$ { at(p) { S1; S2; } }
(iii) Loop-fusion + AT-fusion finish { for (i in D) { async ... = at(D(i)) E1; } } finish { for (i in D) { async ... = at(D(i)) E2; } }	$\Rightarrow$ { finish { for (i in D) { async { t(i)=at(D(i)) {t1=new T(); t1.f1 = E1; t1.f2 = E2; t1;}; ... = t(i).f1; ... = t(i).f2; } } }
[D] Loop-distribution <i>//S1 and S2 are independent of each other and at least one of them contains an at-statement</i> for (i in D) { S1; S2; }	$\Rightarrow$ { for (i in D) { S1; } for (i in D) { S2; } }

Fig. 12. AT-Pruning rules. D: arbitrary distribution; S1, S2: arbitrary statements; E1, E2: arbitrary expressions.

the total number of available places and K be the points over the distribution D (as per the X10 semantics the total points will be $\geq$ available places), and then the total number of place-change operations by un-optimized and optimized code is K and P , respectively. Note that the loop in the input code iterates over the points in D , ordered by the place numbers, and this order is same as the order of the iterations of the generated code. This rule does not require any pre-conditions.

Parallelization of place-change operations. In Figure 12, rules [B(i)] and [B(ii)] parallelizes the invocation of different place-change operations, thereby improving the overall code performance. In rule [B(i)], in the generated code, the place-change operations invoked inside a serial loop in the input loop are parallelized to be evaluated before the main loop. The values thus computed are stored in a temporary rail (array) $r1$ (of size = number of points in D) and reused in the main loop. The dependencies among $S1$, $S2$, and $E1$ are computed using standard techniques [20]. Interestingly, say, “ p ” is the number of places and “ k ” is the size of Distribution D , then the number of remote-communications performed by the code compiled using the synchronization-elimination and place-level strip-mining techniques of Barik et al. [6] are “ $2k$ ” and “ $p + k$,” respectively. In contrast, the rule [B(i)] leads to only “ k ” remote-communications only.

In rule [B(ii)], the parallelization ensures that the place-change operation and serial loop are invoked inside a parallel task. This rule can only be invoked if the loop index variable i and the variable D are not read or modified anywhere after the at-statement.

for(i in D1){ at(D1(i)){ async S1;} S2; }	for(i in D1){ at(D1(i)){ async S1;} for(i in D1){ S2; } }	for(p in place.Places()){ at(p){ for(i in D1(p)){ async S1; } } for (i in D1) { S2; } }	for (p in place.Places()) { async { at(p){ for(i in D1(p)){async S1;} } } for (i in D1) { S2; } }
(a) I/P code	(b) applying rule [D]	(c) applying rule [A]	(d) applying rule [B(ii)]
for(j in D2) { ...=at(D2(j)) ar1(j); ...=at(D2(j)) ar2(j); }	Val r1=new Rail[T](D2.size); Val r2=new Rail[T](D2.size); finish { for(j in D2){ async r1(j)=at(D2(j))ar1(j);} finish { for(j in D2){ async r2(j)=at(D2(j))ar2(j);} } for(j in D2) { ...=r1(j); ...=r2(j); } }	Val r1=new Rail[T](D2.size); Val r2=new Rail[T](D2.size); finish{for (j in D2){async{ t(j)=at(D2(j)){t1=new T(); t1.f1 = ar1(j); t1.f2 = ar2(j); return t1; }; r1(j)=t(j).f1; r2(j)=t(j).f2; } } for(j in D2){...=r1(j);...=r2(j);}	(e) S1 expansion
	(f) applying rule [B(i)]	(g) applying rule [C(iii)]:	

Fig. 13. Example showing the application of the transformation rules of AT-Pruning.

Fusion. In Figure 12, rules [C(i)], [C(ii)], and [C(iii)] show three variants of fusion techniques that we use to reduce the number of place-change operations. Rule [C(i)] is the standard loop fusion transformation applied to codes involving place-change operations. Rule [C(ii)] optimizes consecutive place-change operations (may have been present in the input code or generated as part of some prior transformation like the one using rule [C(i)]) into a single place-change operation. Rule [C(iii)] does both fusion of loops and place-change operations. It creates a temporary object to save the values of E1 and E2. Like rule [C(ii)], rule [C(iii)] is also used to additionally optimize code generated as part of the other translation rules (for example, rule [B(i)]). Even though the rule is shown for only for a sequence involving two `finish` statements (having a specific syntactic pattern), it is applicable to any number of consecutive `finish` statements (with that pattern). The variable `t` is an array of type `T` that can hold the values of E1 and E2. Note that unlike other rules, the rule [C(iii)] does not commute, and hence it is invoked at the end (Figure 11). This rule [C(iii)] reduces the number of place-change operations from “ $n \times K$ ” to K , where “ n ” is the number of such consecutive `finish` statements (with that pattern) and “ K ” is the number of available points in the distribution `D`.

Loop-distribution. The rule [D] is the standard loop-distribution rule, applied only over statements when at least one of them contains an `at`-statement. This rule helps make the code amenable to be optimized using the prior rules ([A], [B], [C]).

Example. Figure 13 shows an example on how the translation rules shown in Figure 12 are applied on an X10 snippet (Figure 13(a)). Figure 13(d) shows the code obtained after applying rules [D], [A], and [B(ii)]. If the statement S1 in Figure 13(a) expands to the code shown in Figure 13(e), then Figure 13(g) shows the translation of S1, after applying rules [B(i)] and [C(iii)]. Note that rule [B(i)] is applied twice on S1 (Figure 13(e)) for the two AT-Expr inside the loop. Because of applying rule [B(i)] twice, the new optimized code became amenable to applying rule [C(iii)]. It can be seen that the input X10 program is exposed to different AT-Pruning rules, shown in Figure 12, until the code reaches a state where no more rules can be applied.

6 AT-OPT AND AT-PRUNING WITH EXCEPTIONS

We now extend AT-Opt and AT-Pruning to handle code that may throw exceptions.

AT-Opt in presence of exceptions. (a) *In the analysis phase:* The object thrown by the `throw` statement at a place p_1 may be caught by the `catch` statement at p_1 or one of its parents in the abstract-place-tree. Considering the complexities in identifying the precise catch statement and its place of execution, we treat the thrown object conservatively and assume that all the fields

AT-Opt	
<pre> t1 = x.f1; t2 = x.f2; ... // create substitute object and initialize fields x1=new X(); x1.f1=t1; x1.f2=t2;... ... </pre>	<pre> ⇒ { Var F:boolean = false; if (x==null) { F = true; t1 = def(..); t2 = def(..); ..} else { t1=x.f1; t2=x.f2; ..} if (F) { x1 =null; } else {x1=new X(); x1.f1=t1; x1.f2=t2;...} ... } </pre>
AT-Pruning	
[B] Parallelization of place-change operations	
<p>(i) AT-Expr parallelization <i>//E1 is independent of S1, S2 and</i> <i>//other instances of E1</i> for (i in D) { S1; <i>//S1 has no at-expression.</i> ... = at(D(i)) E1; S2; }</p>	<pre> ⇒ { Val r1 = new Rail[T](D.size()); Val Ex = new Rail[Exception](D.size()); finish for (i in D) { async { try { r1(i) = at(D(i)) E1; } catch (Exception e) { Ex(i) = e; } } for (i in D) { S1; if (Ex(i)!=null) { throw Ex(i); } else { ... = r1(i); } S2; } } </pre>
[C] Fusion	
<p>(i) Loop-fusion <i>//S1 and S2 are independent of each other</i> <i>//S1 throws no exceptions</i> for(p in place.Places()){at(p) S1;} for(p in place.Places()){at(p) S2;} </p>	<pre> ⇒ { Var Ex:Exception; for (p in place.Places()) { at(p) S1; try { at(p) if(!F) S2; } catch (Exception e) {Ex=e; F=true;}} if(F) throw Ex; Val Ex = new Rail[Exception](D.size()); finish { for (i in D) { async { t(i) = at(D(i)) { t1 = new T(); t1.f1 = E1; try { t1.f2 = E2; } catch (Exception e) {Ex(i) = e; t1; } ... = t(i).f1; ... = t(i).f2; } } } finish for (i in D) { if(Ex(i)) {async throw Ex(i);}} } </pre>
<p>(iii) Loop-fusion + AT-fusion <i>//First for-loop throws no exceptions</i> finish { for (i in D) { async { ... = at(D(i)) E1;}}} finish { for (i in D) { async { ... = at(D(i)) E2;}}} finish for (i in D) { async{try{ r1(i)=at(D(i)) E1;} catch (Exception e){Ex1(i)=e;}}} finish for (i in D) { async{try{ r2(i)=at(D(i)) E2;} catch (Exception e){Ex2(i)=e;}}} }</p>	<pre> ⇒ { finish { for (i in D) { async { t(i) = at(D(i)) { t1 = new T(); try{t1.f1=E1;} catch (Exception e){Ex1(i)=e;} try{t1.f2=E2;} catch (Exception e){Ex2(i)=e;} t1; ... = t(i).f1; ... = t(i).f2; } } } } </pre>
[D] Loop-distribution	
<p><i>//S1 and S2 are independent of each other</i> <i>//One of S1 and S2 contains an at-statement</i> <i>//S2 throws no exceptions</i> for (i in D) { S1; S2; }</p>	<pre> ⇒ { Val Ex = new Rail[Exception](D.size()); for (i in D) { try { S1; } catch (Exception e){Ex(i) = e;break;}} for (i in D) if (Ex(i)==null) S2; else throw Ex(i); } </pre>

Fig. 14. AT-Opt and AT-Pruning in presence of Exceptions.

reachable from that object are read in the throw statement. Similarly, the argument of each catch block is also treated conservatively. Considering that the exceptions are rarely thrown, our chosen conservative design does not reduce our gains much. (b) *In the code generation phase:* Before an at-construct, we emit code (of the form, $t = x.f$) that eagerly dereference the object fields to copy their values into temporaries. If the variable x points-to *null*, then such a dereference will throw a `NullPointerException` earlier than the original dereference point (inside the at-construct) in the input program. Note: No other exception may be thrown because of our generated code. To preserve the semantics of the generated code, as shown in Figure 14, instead of the simple codes (shown in the left) we emit code shown in the right. Here, we first check if the variable points-to *null*, and if so, we set a flag F to true and initialize the temporaries $t1$, $t2$, and

so on, to their default initial values. Later, inside the `at`-construct, we create a substitute object only if F is false.

AT-Pruning in presence of exceptions. Except for the rules [A], [B(ii)], and [C(ii)], the other rules in Figure 12 may alter the execution order of the statements and hence need to be handled in a special manner in the presence of exceptions. Figure 14 presents the modified rules.

Rule [B(i)]: If any exception is thrown during the evaluation of $E1$, then we save the exceptions in an array and throw the first exception in the main loop (after executing $S1$). **Rule [C(i)]:** If an exception is thrown in any instances of $S2$, then (a) the exception is saved, (b) the latter instances of $S2$ are not executed, and (c) the saved exception is thrown outside the loop. **Rule [C(iii)]:** We catch the exceptions thrown in the main loop and throw them outside the loop. We also do Loop- and AT-fusion in presence of try-catch statements (to handle the output of rule [B(i)]). **Rule [D]:** We catch any exception thrown, while executing the instances of $S1$ and execute only those instances of $S2$ whose corresponding instance of $S1$ executed and did not throw any exception. The exceptions thrown from $S1$ are captured in an array Ex and thrown before executing the corresponding $S2$.

7 DISCUSSION

We now discuss some interesting underlying points about our proposed optimization scheme.

(a) **AT-Com in the compiler.** Both AT-Opt and AT-Pruning are high-level optimizations that do not interfere negatively with other high-level optimizations. AT-Opt requires a pre-pass of expression simplification (to three-address code). AT-Opt and AT-Pruning should be invoked before any other high-level-optimization that may change the structure of `at`-constructs. AT-Opt and AT-Pruning can be applied in any order to get the combined improvements.

(b) **Ambiguous object.** Our idea of ambiguous objects helps classify objects that *need* to be copied fully (non-ambiguous) and those that have to be *conservatively* copied fully (ambiguous). We believe that such a classification is novel and can enable future optimizations.

(c) **Partitioning Aggregate Object fields.** Access to object field of type vectors, arrays, and rails are handled using an element insensitive approach. The reason for not using element sensitive approach is because of more complex analysis and additional data structures are required to keep track of each element of such data types. So (i) AT-Opt treats vectors, arrays, and rails, along with all their constituent elements as a single object. Thus read/write to an element (or a sub-partition) is considered as read/write of that object. (ii) Further, each write to an element (or a sub-partition) of data type vectors, arrays, and rails is also considered as a read to those objects. In our experience, we have found that programs rarely access non-distributed array elements. Another possible way to handle this complex problem of partitioning aggregate data is by asking the programmer to specify the sub-regions that are shared; such an approach is used by Legion [7].

(d) **Transient and GlobalRef fields.** We ignore the fields declared transient (not required to be copied) or GlobalRef (no scope to optimize).

(e) **Code generation for substitute object.** During code generation (Section 3.2), AT-Opt emits code to create a new substitute object for different objects and initializes its fields from the temporaries created in the previous phase (Figure 8(a)). We can further optimize this part by avoiding the creation of new substitute objects (altogether) and replacing the corresponding field dereferences with the temporaries. This optimization can be done only if the object under consideration is not passed to any function call (including as the `this` argument). Note: A dummy constructor is added for each user-defined type of the input X10 program to assist in creating the substitute objects.

(f) **Code generation for the statement $a=b$.** During code generation, if the object o_i pointed-to by b is non-ambiguous, and the code generator has not emitted code to create substitute object for o_i , then it indicates that o_i is not dereferenced/used in the body of the `at`-construct; likely dead-code. Hence, we eliminate the statement altogether.

```

public class mst {
    val START = 0;
    val JOIN = 1;
    val CHANGE = 2;
    var adj_graph:Array[Long];
    var nodes:Long;
    val Infinity = Long.MAX_VALUE;
    var nSet:DistArray[node];
    var loadValue:long=0;
    var nval:DistArray[long];
    var R: Region; var D: Dist;
    var cagain:DistArray[boolean];
    var tagain:DistArray[boolean];
    var invagain:DistArray[boolean];
    var setCheck:DistArray[boolean];
    var checkflag:DistArray[boolean];
    ...
}

def setChildSignal():boolean {
    finish {
        for(i in this.D) {
            async at(this.D(i)) { //pco#1
                this.nSet(i).ss=this.nSet(i).tss;
                if(!(this.loadValue.equals(0))) {
                    val pt:Long = i(0);
                    this.nval(i)=this.weight(this.nval(i)+pt);
                } } } /* finish */

    var retVal:boolean = false;
    for(i in this.D) {
        //pco#2
        var atVal:boolean=at(D(i)) this.setCheck(i);
        if(atVal) retVal = true;
    }
    return retVal; } /*setChildSignal*/ } /*mst*/

```

Fig. 15. IMSuite kernel: MST, data members and member functions.

(g) **Remote reads/writes instead of AT-Opt.** The alternative of using remote references for local reads/writes of non-distributed objects is unsuitable, as the writes to “local” memory of a place will be visible to other tasks running on other places, which violates the underlying PGAS semantics.

(h) **Handling of at-expressions by AT-Opt.** They are handled in the same way as the at-statements.

(i) **Impact of AT-Opt.** Note that if all the fields of the reachable objects are scalars or distributed arrays, or most of the data accessible from shared objects are accessed inside the at-construct, then the gains due to AT-Opt compared to X10 compiler would be minimal.

(j) **Why flow-sensitive analysis?** AT-Opt uses flow-sensitive analysis to precisely identify object fields dereferenced across remote places. In our running example, for the place-change operation at line 4, the flow-insensitive approach will copy all fields (r1 and r2) of the object pointed by a, though only r1 is required during program execution.

(k) **Profitability of the rules of AT-Pruning.** The transformation rules discussed as part of AT-Pruning are invoked syntactically, and profitability is not explicitly evaluated. That is, a rule of the form $X \Rightarrow Y$ is invoked every time we encounter some code of the form X . Consequently, one may argue that some of the invocations may be redundant and even lead to performance deterioration. For example, if Rule [D] (Figure 12) is invoked first, but it led to no actual optimization (say, because of issues related to dependencies), then the translation is redundant and in some cases leads to inefficient codes (executing the loop headers twice). However, in our evaluation, we found that such cases did not lead to any visible performance deterioration.

(l) **Comparison against prior work.** Figure 15 shows the code of the MST kernel from IMSuite. For brevity, we have shown all 16 data members and only the setChildSignal method of the MST kernel. The setChildSignal method contains two place-change operations (pco#1 and pco#2) to check if any child can start processing in parallel. Figure 16(a) and 16(b) shows the optimized MST kernel generated by our proposed AT-Com and Barik et al. [6] work, respectively. For pco#1: AT-Com applies rule [A] of Figure 12 (AT-Pruning), and then AT-Opt precisely identifies the object fields (nset, nval, loadValue, and D) dereferenced at target places and emits the optimized code such that only required fields are copied during program execution, whereas Barik et al. [6] use a simple scheme of scalar replacement to reduce the amount of data communicated across places. In their

```

def setChildSignal():boolean {
  finish {
    for(p in place.Places()) {
      val t1=this.nSet; val t3=this.nval;
      val t2=this.loadValue; val t4=this.D;
      at(p) {
        val t5:mst = new mst();
        t5.nSet=t1; t5.nval=t3;
        t5.loadValue=t2; t5.D=t4;
        for(i in t5.D(p)) { async {
          t5.nSet(i).ss=t5.nSet(i).tss;
          if(! (t5.loadValue.equals(0))) {
            val pt:Long = i(0);
            t5.nval(i)=t5.weight(t5.nval(i)+pt);
          } } } } } /* finish */

    var retVal:boolean = false;
    val r1=new Rail[boolean](this.D.size());
    finish for(i in this.D) {
      async r1(i)=at(D(i)) this.setCheck(i);
    }
    for(i in this.D) {
      var atVal:boolean=r1(i);
      if(atVal) retVal = true;
    }
    return retVal;
  }
}

```

(a) AT-Com optimized code.

```

def setChildSignal():boolean {
  finish {
    for(i in this.D) {
      val t6 = this.loadValue;
      async at(this.D(i)) {
        this.nSet(i).ss=this.nSet(i).tss;
        if(! (t6.equals(0))) {
          val pt:Long = i(0);
          this.nval(i)=this.weight(this.nval
            (i)+pt);
        } } } /* finish */

    var retVal:boolean = false;
    for(p in place.Places()) {
      val sd = this.D(p);
      val t7=new Rail[boolean](sd.size());
      finish at(p) async {
        var j:long = 0;
        for(i in sd) {
          val ind=j++; val v=this.setCheck(i);
          at(t7) async t7(ind)=v;
        }
        var j:long = 0;
        for(i in sd) {
          if(t7(j++))
            retVal = true;
        } /* for p in place.Places */
        return retVal;
      }
    }
  }
}

```

(b) Code optimized by Barik et al. [6].

Fig. 16. Optimized MST kernel by (a) AT-Com and (b) Barik et al. [6]. The changed code is shown in red color.

optimized code, the scalar field `loadValue` is saved to `t6` and its occurrence is replaced by `t6`. But their optimized code will still copy the complete `this` object, because (i) `nset` and `nval` are non-scalar fields (note that Barik et al. [6] scalar replacement for array access cannot be applied on distributed arrays, because their elements are owned by target places), and (ii) an access of the `this` object during the function call `this.weight`.

For `pco#2`: AT-Com applies rule [B(i)] of Figure 12, whereas synchronization-elimination and place-level strip-mining rules are applied by Barik et al. [6]. Though both the approaches execute the place-change operations in parallel, our technique leads to faster code (detailed in Section 5). (m) **Time and Space complexities**: The worst-case time-complexity of AT-Opt is bound by that of points-to analysis $O(N^3)$, where N is the size of the program. Similarly, the worst-case space-complexity is bound by the size of the points-to graph $O(N^2)$.

In case of AT-Pruning, the number of times the rules [A]–[D] in Figure 12 can be applied on the input program is $O(N)$, and similarly these rules may add $O(N)$ new statements, which can lead to the invocation of the above rules at most $O(N)$ number of times. Each rule may take $O(N)$ time and hence the worst-case complexity of AT-Pruning is $O(N^2)$; this is assuming that the dependence analysis has been pre-computed and incrementally updated after each rule invocation. AT-Pruning does not use any auxiliary memory for its processing. Alternatively, we can invoke

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Name	I/P	#at-constructs				serialized-data (GB)				AT-Com		AT-Opt		AT-Pruning	
		Base		AT-Com		Base	AT-Com	AT-Opt	AT-Pruning	% reduction		% reduction		% reduction	
		Stat	Dyn	Stat	Dyn					c-m(I)	c-m(A)	c-m(I)	c-m(A)	c-m(I)	c-m(A)
BF	256	45	6,403	42	4,618	3.00	0.09	0.12	2.13	95.83	51.90	91.94	33.48	19.80	23.18
DST	256	101	15,660	93	7,755	7.39	0.40	0.50	3.53	96.76	64.73	88.58	31.23	35.96	46.07
BY	128	69	550,254	65	539,332	68.08	0.25	0.25	66.73	91.54	28.65	91.14	25.92	1.16	0.02
DR	256	39	489,944	37	489,434	120.07	0.38	0.39	119.70	96.02	40.05	95.80	39.42	0.02	0.39
DS	256	195	542,730	184	514,680	140.09	0.24	0.26	126.37	88.21	53.51	86.98	35.06	4.83	28.54
KC	256	121	84,178	110	72,449	0.32	0.15	0.16	0.30	34.42	22.01	21.97	7.75	22.53	15.99
DP	256	97	68,651	92	57,176	32.18	0.15	0.15	26.55	96.72	49.79	95.60	38.08	9.98	14.07
HS	256	130	400,477	124	7,777	1.01	0.01	0.22	0.01	92.75	90.87	22.05	5.56	91.98	90.85
LCR	256	48	197,639	45	66,824	0.47	0.08	0.09	0.16	86.05	84.46	22.75	8.61	80.61	82.77
MIS	256	68	18,814	62	15,244	8.94	0.13	0.13	7.19	96.67	52.19	88.55	42.78	16.22	21.09
MST	256	254	193,103	238	118,643	76.48	8.32	10.32	35.44	95.95	62.70	78.95	25.90	47.20	41.95
VC	256	86	5,126	80	2,066	2.27	0.10	0.12	0.76	97.07	65.05	86.39	30.68	59.78	37.05

Fig. 17. Characteristics of the IMSuite kernels. Abbreviations: c-m: cache misses; (I) Intel; (A) AMD.

the complete dependence analysis (if the dependencies are not incrementally computed) before rules [B(i)], [C(i)], and [D] to compute dependencies between statements and the complexity of the dependence analysis will have to factored in.

Overall: We find that the time-complexity of AT-Com is $O(N^3)$ and the space-complexity is $O(N^2)$, where N is the program size.

8 IMPLEMENTATION AND EVALUATION

In this section, we evaluate our proposed optimization scheme AT-Com on two different systems: (i) a four-node Intel system, where each node has two Intel E5-2670 2.6GHz processors, 16 cores per processor, 64GB RAM per node, and 20MB cache-per core and (ii) a two-node AMD system, where each node has an AMD Abu Dhabi 6376 processor, 16 cores per processor, 512GB RAM per node, and 2MB cache per core. We implemented AT-Opt and AT-Pruning, the constituent optimizations of AT-Com, in the x10v2.6.0 compiler x10c++ (C++ backend). Based on the ideas from the insightful paper of Georges et al. [12], we report the execution times by taking a geomean over 30 runs.

We evaluated AT-Com using 12 benchmark kernels from IMSuite [14]: breadth first search (BF computes the distance of every node from the root and DST computes the BFS tree), byzantine consensus (BY), routing table creation (DR), dominating set (DS), maximal independent set (MIS), committee creation (KC), leader election (DP, for general network; HS, for bidirectional ring network; and LCR, for unidirectional ring network), spanning tree (MST), and vertex coloring (VC). We also studied many other benchmarks made available in the X10 distribution, but none of them met our selection requirements: (a) presence of at-construct in the program and (b) de-reference of object fields at a remote place (say, p1), such that the object is not created in the remote place p1.

Figure 17, columns 2 to 4, shows the chosen input sizes and the number of remote-communications (number of at statements) during both compile-time and runtime for the chosen input. For these benchmark kernels, we found that the maximum height of Abstract-place-tree was at most three. For all the benchmarks kernels, the chosen input size was the largest input such that on our 32-core Intel system, when the input program, compiled using the default compiler, is run by setting X10_NPLACES=2, it does not take more than an hour to execute and does not run out of memory. We executed the chosen kernels on the specified inputs by varying the number of places (in powers of two) and threads per place such that at any point of time the total number of threads ($= \#places \times \#num-threads-per-place$) is equal to the number of cores. This is achieved by setting

(a) Intel system speedups; totalCores=128. Speedup=(exec-time using Base / exec-time using AT-Com).

(b) AMD system speedups; totalCores=32. Speedup=(exec-time using Base / exec-time using AT-Com).

Fig. 18. Speedups for varying number of places (#P) and threads (#T). Config $C_i \equiv \#P = i$ and $\#T = \text{totalCores}/i$.

the runtime environment variable `X10_NPLACES` and `X10_NTHREADS` (threads per place) appropriately. The default X10 runtime divides the places equally among all the provided hardware nodes.

8.1 Impact of AT-Com

In this section, we present the evaluation of the overall impact of AT-Com (uses both AT-Opt and AT-Pruning) by comparing against Base, the baseline version without any communication optimizations. In Figure 17, columns 5 and 6 show the number of static and dynamic counts of the at-constructs in the AT-Com optimized code. It shows that AT-Com leads to a significant reduction in the number of static and dynamic at-constructs: 6.5% and 26%, respectively. Naturally, these reductions are obtained only from the AT-Pruning optimization (AT-Opt does not change the number of at-constructs). In Figure 17, columns 7 and 8 report the amount of data (excluding some common metadata and body of the at-construct) serialized during the execution of kernel programs, in the context of Base and AT-Com, respectively. As it can be seen, compared to Base the AT-Com optimized code leads to a large reduction in the amount of serialized data (6× to 583×). Note that the amount of data serialized is independent of the number of places and is only dependent on the number of at-constructs and the data serialized at each of them.

Figure 18(a) and Figure 18(b) show the speedups achieved by using AT-Com on the four-node Intel system and AMD system, respectively, for varying number of places and threads. It can be seen that with respect to Base, AT-Com optimizer achieved significant speedups: geomean of 18.72× on the Intel system and 17.83× on the AMD system. It can be seen that the speedups for all benchmarks across the Intel and AMD systems are consistent, with the reduction in the number of

Fig. 19. Individual impact of AT-Opt (say X) and AT-Pruning (say Y), and the combined impact (AT-Com-(X+Y)).

at-constructs and the reduction in the amount of serialized data. For example, for the KC kernel, the relatively lower gains (though significant) is due to the low reduction in the number of dynamic at-constructs and the amount of serialized data. The exact speedup may vary, depending on the input program, input, and hardware (including the available cores, memory, cache and so on).

Compared to the rest of the kernels showing very high speedups, the speedups of kernels BY, HS, and LCR look lower. This is because for these three kernels, the relative amount of time they spend in computation (compared to the amount of serialized data and number of place-change operations) is much higher (in contrast to kernels like BF, DR, DS, and so on).

8.2 Impact of individual components of AT-Com

We have also studied the individual impacts of AT-Opt and AT-Pruning on the overall gains resulted due to AT-Com. Figure 19 shows the individual contributions of AT-Opt and AT-Pruning to the overall speedups and the additional impact realized by the combination (AT-Opt+AT-Pruning). On average, we see that AT-Opt and AT-Pruning individually contribute 40.5% and 11% of the overall gains. The rest of the gains are realized because of the combination AT-Opt + AT-Pruning. This combined effect was very high especially for BF, DST, DR, MIS, MST, and VC.

Contributions of AT-Opt. Figure 17, column 9, reports the amount of data serialized during the execution of kernel programs compiled by AT-Opt. Compared to Base, the AT-Opt optimized code leads to a reasonable reduction in the amount of serialized data ($2\times$ to $527\times$) and achieved large speedups: geomean of $9.30\times$ on the four-node Intel system and $5.57\times$ on the two-node AMD system; the detailed breakup is skipped for space.

In Figure 19, it can be seen that except for KC, HS, and LCR, the contribution to speedups by AT-Opt across the Intel and AMD systems is consistently high. For kernels KC, HS, and LCR, the contributions to speedups are not substantial. As discussed in the previous section, this is due to the amount of data getting communicated across places (in Base, itself) is very less; consequently, the reduction in the communicated data is also less (order of few hundred MBs; see Figure 17). For the rest of the benchmarks, AT-Opt leads to significant amount of gains in the execution time (in line with the reduction in the communicated data).

Contributions of AT-Pruning. Figure 17, columns 5 and 6, reports the total number of remote-communications (place-calls) made during the execution of these kernel programs, in the context of AT-Pruning. As it can be seen, compared to Base the AT-Pruning transformed code leads to the reduction in remote-communication (place-change operations) across places ($1.01\times$ to $51.5\times$) and also results in the reduction in the amount of data serialized ($1.01\times$ to $101\times$; column 10, Figure 17). Figure 20(a) and Figure 20(b) show the speedups achieved by using AT-Pruning, on the Intel and AMD systems, respectively, for varying number of places and threads. We see that

(a) Intel system speedups; totalCores=128. Speedup=(exec-time using Base / exec-time using AT-Pruning).

(b) AMD system speedups; totalCores=32. Speedup=(exec-time using Base / exec-time using AT-Pruning).

Fig. 20. Speedups for varying number of places (#P) and threads (#T). Config $C_i \equiv \#P = i$ and $\#T = \text{totalCores}/i$.

compared to Base, the AT-Pruning optimizer achieved reasonable speedups: geomean $1.99\times$ (Intel) and $2.17\times$ (AMD).

Figures 19 and 20 show that except for BY and DR the contribution to speedups by AT-Pruning across both the systems are reasonable. The speedup depends on the maximum possible number of instances where AT-Pruning can be applied. The speedups are directly proportional to the reduction and parallel execution of place-change operations. For kernels BY and DR, the speedups are not substantial. This is because in these kernels, the impact of AT-Pruning is quite low (columns 5 and 6). For HS and LCR, AT-Pruning leads to a large reduction in the place-change operations ($>100K$ number of place-change operations each). This in turn reduces the amount of serialized data considerably and leads to high speedups. For the rest of the kernels, AT-Pruning led to lesser reduction of place-change operations and this is reflected in the resulting speedups.

8.3 Evaluation on a Single-node (Shared Memory Environment) Setup

Figure 21 shows the geomean speedups achieved by AT-Com on an Intel (32 cores) and an AMD (16 cores) single node system. We can see that AT-Com leads to significantly high speedups, even on a single node system: geomean $7.13\times$ and $9.94\times$ on the Intel and AMD system, respectively. The corresponding numbers for AT-Opt and AT-Pruning are $3.06\times$, $3.44\times$ and $2.20\times$, $2.38\times$, respectively. Naturally, the speedups on the distributed-memory system are more because of the reduction in inter-node communication.

To understand the sources of the obtained speedups, we studied the cache behavior of the generated codes on the single node Intel and AMD systems. Figure 17, columns 11–16, reports the

Fig. 21. GeoMean speedups for varying number of places (#P) and threads (#T) for single node Intel (#cores = 32) and AMD (#cores = 16). Config: $C_i \equiv \#P = i$ and $\#T = \#cores/i$.

geomean % reduction (compared to Base) in cache-misses (for AT-Com 34–97% on the Intel system and 22–90% on the AMD system) across C_2 , C_4 , C_8 , C_{16} , and C_{32} . Thus, the overall gains = gains from reduced copying and data-transfer + reduced memory-access cost due to reduced cache-misses. It can be seen that reduction in cache-misses is less on the AMD system. We conjecture the reason to be related to the cache size: Since AT-Com reduces the amount of memory usage, which in turn leads to less cache pollution and its impact is more visible on the Intel system with larger cache.

8.4 Impact of the Design Choices

We now discuss the empirical impact of two main design choices made as part of AT-Opt.

Inter-procedural analysis. For the benchmarks under consideration, we found that the inter-procedural component was essential for getting the reported large speedups. This is because the intra-procedural analysis alone could impact only two benchmarks (BY and DP). Even there, the impacted at-constructs were not part of the main computation and hence led to no gains.

Conservative handling of ambiguous-objects. In the chosen kernels that met our selection requirements (Section 8), the impact due to ambiguous objects was non-existent. This is not unexpected, because in HPC codes, it is not common to conditionally create objects and dereference them at remote places. However, we still show the procedure to deal with them to make sure that our proposed technique is sound (see Figure 7 and the discussion thereof).

8.5 Comparison to the Prior Work

Barik et al. [6] use a simple scheme of scalar replacement to reduce the amount of data communicated across places; see Section 7 for a illustrative comparison. While that scheme identifies a subset of opportunities identified by our proposed technique and can be effective in some cases, for the IMSuite benchmarks, the scalar-replacement scheme had no impact whatsoever. This is because in these benchmarks the objects, whose field accesses are optimized by our technique, are passed as arguments (receiver or parameters) to functions, and, consequently, scalar replacement is not performed, as it may require non-trivial modifications to the method signatures. Further, many of those were non-scalar fields (violates the requirement to perform scalar replacement [6]). This renders a comparison of AT-Opt against the scalar-replacement scheme of Barik et al. redundant. Similarly, the impact of rule [B(i)] of Figure 12 and its counterpart rules of Barik et al. (synchronization-elimination and place-level strip-mining) was minimal and thus not detailed.

8.6 Overall Summary of the Evaluation

We have done a detailed evaluation and found that the actual speedup varies depending on multiple factors: (1) Number of executed at-constructs. (2) Amount of data getting serialized during each communication. (3) Amount of other components of remote communication (metadata such as

runtime-type information, data related to the body of the `at-construct`, and so on). (4) Time taken to perform inter-place communication. (5) The nature of the input, runtime/OS related factors and the hardware characteristics. The factor (1) is impacted by AT-Pruning and factor (2) is the only one that is different between Base and AT-Opt optimized codes. However factor (1) impacts factor (2) and the impact of factor (2) can be felt on (4) as well. Since AT-Com reduces the factors (1) and (2) (consequently, factor (4)), it leads to significant performance gains.

9 RELATED WORK

There have been many prior works [3, 5, 6, 10] that aim to reduce the communication overheads across places resulting from redundant data transfers. Barik and Sarkar [5] eliminate memory loads by scalar replacement in single place X10 programs. The scalar-replacement scheme of Barik et al. [6] targets multi-place X10 programs and has similarities with AT-Opt but with the following differences: (i) They handle only scalar fields; AT-Opt goes beyond that and reduces remote-data transfers involving heap objects. (ii) They do not handle writes to fields; AT-Opt can handle reads and writes of both mutable and immutable-fields. (iii) They cannot handle “ambiguous” objects (Section 3.2, Figure 7); AT-Opt can. (iv) Importantly, unlike AT-Opt, their scheme cannot be applied where the object (whose fields may have to scalar-replaced) is passed to a function as an argument/receiver. Because of these points, the impact of their scalar-replacement scheme was negligible on the IMSuite kernels (see Section 8).

Besides the scalar-replacement and program transformation techniques, Barik et al. [6] also present other compiler optimizations to reduce communication and synchronization overheads: task localization, object splitting for array of objects, replicating arrays across places, distributing loops to specialize local-place accesses, and so on. We believe that these techniques can be effectively used along with our techniques in an orthogonal manner by invoking AT-Com first and then these optimizations (that may change the structure of loops/`at-construct`).

There have been prior works [3, 10] that aim to optimize communication of fine-grain data by eliminating redundant communication, use of split-phase communication, and coalescing. Similarly, Hiranandani et al. [16, 17] have developed a framework called Fortran D that reduces communication overheads by applying optimizations like message vectorization, message coalescing, message aggregation, and pipelining. These techniques are further extended by Kandemir et al. [18] to optimize the global communication. Our proposed work targets general communication (not just fine-grain communication) and can be invoked before their schemes to take advantage of both.

Sanz et al. [25] optimize the communication routines and block and cyclic distribution modules in Chapel [8] by performing aggregation for array assignments. Paudel et al. [21] propose a new coherence protocol in the X10 runtime to manage mostly read shared variables. Our proposed technique can be used on top of such runtime optimizations to further improve the performance.

Sharma et al. [27] perform affine loop optimization using modulo unrolling in Chapel [8] to access elements of array distributed across different locale (equivalent to places in X10). They propose a technique named zipper iteration that uses modulo unrolling to unroll loops and then gathers array elements from remote locale to a local buffer via one bulk get message and assign it to the current locale. Our rule [B(i)] of Figure 12 can be used on top of it to further improve the performance. Similarly, Pellegrini et al. [22] perform precise data dependence analysis with the help of polyhedral model to restore the compiler analysis techniques to remove the communication overlap in MPI programs.

There have been many works on points-to and shape analysis [4, 11, 13, 23, 24, 29]. Chandra et al. [9] use a dependent-type system to reason about the locality of X10 objects. We extend the escapes-to-connection graph of Agarwal et al. [1] to reason about places and their accessed objects.

10 CONCLUSION

In this article, we present a novel scheme AT-Com to reduce the communication overheads by paying close attention to objects getting copied across *places* and number of place-change operations in X10 programs. We have implemented AT-Opt and AT-Pruning, the constituent optimizations of AT-Com, in the x10v2.6.0 compiler and evaluated the performance on two different systems (a four-node $\times$ 32-core Intel system and a two-node $\times$ 16-core AMD system). We show that AT-Com leads to significant gains in execution time with speedups of 18.72 $\times$ and 17.83 $\times$ on the Intel and AMD systems, respectively. Additionally, the experimental results show that the AT-Com optimized programs scale better than the baseline versions. Though we discussed AT-Com in the context of X10, we believe that it can be applied to other PGAS languages like Chapel, HJ, and so on.

REFERENCES

- [1] S. Agarwal, R. Barik, V. K. Nandivada, R. K. Shyamasundar, and P. Varma. 2008. Static detection of place locality and elimination of runtime checks. In *APLAS*, G. Ramalingam (Ed.). San Francisco, CA, 53–74.
- [2] S. Agarwal, R. Barik, V. Sarkar, and R. K. Shyamasundar. March 2007. May-happen-in-parallel analysis of X10 programs. In *PPoPP*. 183–193. DOI : <https://doi.org/10.1145/1229428.1229471>
- [3] M. Alvanos, M. Farreras, E. Tiotto, J. N. Amaral, and X. Martorell. 2013. Improving communication in PGAS environments: Static and dynamic coalescing in UPC. In *ICS*. 129–138. DOI : <https://doi.org/10.1145/2464996.2465006>
- [4] L. O. Andersen. 1994. *Program analysis and specialization for the C programming language*. Ph.D. Dissertation. University of Copenhagen.
- [5] R. Barik and V. Sarkar. September, 2009. Interprocedural load elimination for dynamic optimization of parallel programs. In *PACT*. 41–52. DOI : <https://doi.org/10.1109/PACT.2009.32>
- [6] R. Barik, J. Zhao, D. Grove, I. Peshansky, Z. Budimlic, and V. Sarkar. 2011. Communication optimizations for distributed-memory X10 programs. In *IPDPS*. 1101–1113. DOI : <https://doi.org/10.1109/IPDPS.2011.105>
- [7] M. Bauer, S. Treichler, E. Slaughter, and A. Aiken. 2012. Legion: Expressing locality and independence with logical regions. In *SC*. IEEE Computer Society Press, Los Alamitos, CA, Article 66, 11 pages. <http://dl.acm.org/citation.cfm?id=2388996.2389086>.
- [8] B. L. Chamberlain, D. Callahan, and H. P. Zima. 2007. Parallel programmability and the chapel language. *Int. J. High Perform. Comput. Appl.* 21, 3 (Aug. 2007), 291–312. DOI : <https://doi.org/10.1177/1094342007078442>
- [9] S. Chandra, V. Saraswat, V. Sarkar, and R. Bodik. 2008. Type inference for locality analysis of distributed data structures. In *PPoPP*. 11–22. DOI : <https://doi.org/10.1145/1345206.1345211>
- [10] W. Chen, C. Iancu, and K. Yelick. 2005. Communication optimizations for fine-grained UPC applications. In *PACT*. 267–278. DOI : <https://doi.org/10.1109/PACT.2005.13>
- [11] J. Choi, M. Gupta, M. Serrano, V. C. Sreedhar, and S. Midkiff. Nov, 1999. Escape analysis for Java. In *OOPSLA*. 1–19. DOI : <https://doi.org/10.1145/320384.320386>
- [12] A. Georges, D. Buytaert, and L. Eeckhout. Oct, 2007. Statistically rigorous Java performance evaluation. In *OOPSLA*. 57–76. DOI : <https://doi.org/10.1145/1297027.1297033>
- [13] R. Ghiya and L. J. Hendren. 1996. Is it a tree, a DAG, or a cyclic graph? A shape analysis for heap-directed pointers in C. In *POPL*. 1–15. DOI : <https://doi.org/10.1145/237721.237724>
- [14] S. Gupta and V. K. Nandivada. 2015. IMSuite: A benchmark suite for simulating distributed algorithms. *J. Parallel Distrib. Comput.* 75 (2015), 1–19. DOI : <https://doi.org/10.1016/j.jpdc.2014.10.010>
- [15] Habanero. 2009. Habanero Java. Retrieved from <http://habanero.rice.edu/hj>.
- [16] S. Hiranandani, K. Kennedy, and C. Tseng. 1992. Compiling fortran D for MIMD distributed-memory machines. *Commun. ACM* 35, 8 (Aug. 1992), 66–80. DOI : <https://doi.org/10.1145/135226.135230>
- [17] S. Hiranandani, K. Kennedy, and C. Tseng. November, 1991. Compiler optimizations for fortran D on MIMD distributed-memory machines. In *SC*. 86–100. DOI : <https://doi.org/10.1145/125826.125886>
- [18] M. Kandemir, P. Banerjee, A. Choudhary, J. Ramanujam, and N. Shenoy. 1999. A global communication optimization technique based on data-flow analysis and linear algebra. *ACM Trans. Program. Lang. Syst.* 21, 6 (Nov. 1999), 1251–1297. DOI : <https://doi.org/10.1145/330643.330647>
- [19] S. S. Muchnick. 1997. *Advanced Compiler Design and Implementation*. Morgan Kaufmann Publishers Inc., USA.
- [20] V. K. Nandivada, J. Shirako, J. Zhao, and V. Sarkar. 2013. A transformation framework for optimizing task-parallel programs. *ACM Trans. Program. Lang. Syst.* 35, 1 (2013), 1–48.
- [21] J. Paudel, O. Tardieu, and J. N. Amarai. 2014. Optimizing shared data accesses in distributed-memory X10 systems. In *HiPC*. 1–10. DOI : <https://doi.org/10.1109/HiPC.2014.7116889>

- [22] S. Pellegrini, T. Hoefler, and T. Fahringer. 2012. Exact dependence analysis for increased communication overlap. In *Recent Advances in the Message Passing Interface*, J. L. Träff, S. Benkner, and J. J. Dongarra (Eds.). 89–99.
- [23] N. Rinetzky and S. Sagiv. 2001. Interprocedural shape analysis for recursive programs. In *CC*. 133–149. <http://dl.acm.org/citation.cfm?id=647477.727768>.
- [24] A. Salcianu and M. Rinard. 2001. Pointer and escape analysis for multithreaded programs. In *PPoPP*. 12–23. DOI: <https://doi.org/10.1145/379539.379553>
- [25] A. Sanz, R. Asenjo, J. Lopez, R. Larrosa, A. Navarro, V. Litvinov, S. Choi, and B. L. Chamberlain. 2012. Global data re-allocation via communication aggregation in chapel. In *SBAC-PAD*. 235–242. DOI: <https://doi.org/10.1109/SBAC-PAD.2012.18>
- [26] V. Saraswat, B. Bloom, I. Peshansky, O. Tardieu, and D. Grove. 2016. X10 Language Specification Version 2.6.0. Retrieved from <http://x10.sourceforge.net/documentation/languagespec/x10-260.pdf>.
- [27] A. Sharma, D. Smith, J. Koehler, R. Barua, and M. Ferguson. 2014. Affine loop optimization based on modulo unrolling in chapel. In *PGAS*. Article 13, 12 pages. DOI: <https://doi.org/10.1145/2676870.2676877>
- [28] A. Thangamani and V. K. Nandivada. Nov, 2018. Optimizing remote data transfers in X10. In *Proceedings of the 27th International Conference on Parallel Architectures and Compilation Techniques, PACT 2018*. 1–15.
- [29] J Whaley and M Rinard. Nov, 1999. Compositional pointer and escape analysis for Java programs. In *OOPSLA*. 187–206. DOI: <https://doi.org/10.1145/320384.320400>

Received March 2019; revised May 2019; accepted July 2019