eskimo: EXPERIMENTING SKELETONS ON THE SHARED ADDRESS MODEL

MARCO ALDINUCCI
Computer Science Dept. – University of Pisa
Via Buonarroti 2, I–56127 PISA, Italy
(aldinuc@di.unipi.it)

ABSTRACT

We discuss the lack of expressivity in some skeleton-based parallel programming frameworks. The problem is further exacerbated when approaching irregular problems and dealing with dynamic data structures. Shared memory programming has been argued to have substantial ease of programming advantages for this class of problems. We present eskimo library which represents an attempt to merge the two programming models by introducing skeletons in a shared memory framework.

Keywords: Skeletons, dynamic data structures, software DSM, cluster of workstations

1. Introduction

The development of efficient parallel programs is a quite hard task. Besides coding the algorithm, the programmer must also take care of the details involved in parallelism exploitation, i.e. concurrent activity set up, mapping and scheduling, synchronization handling and data allocation. In unstructured, low level parallel programming approaches these activities are usually fully in charge of the programmer and constitute a difficult error prone programming effort. From Cole's seminal work [1] skeleton research community has been active in experimenting new technologies in order to simplify parallel programming by raising the level of abstraction.

In the past decade we designed and developed several skeleton-based parallel programming environments and we tested their effectivity on a number of real world applications. Even if the skeletal approach has been proved to be effective for some of them, the overall feedback we received cannot be considered fully satisfactory. Actually a lack of expressivity emerged, at least for some complex applications.

In this paper we present eskimo [Easy SKeleton Interface (Memory Oriented)] a new skeletal programming environment which represents a preliminary attempt to defeat expressivity lacks emerged in skeletal languages, especially approaching irregular problems and dealing with dynamic data structures. eskimo is based on shared address programming model; its run-time is built upon a software DSM.

In the next section we present a brief (self-critical) history of parallel programming frameworks evolution. In section *3* we present eskimo design principles. In section *4* we discuss the pay-back we expect from the skeletal approach. The paper is completed by some experimental results (sec. *5*) and the related work (sec. *6*).

2. Motivation and Historical Perspective

Historically a couple of works are due particular attention: the P^{3}L project [2] and the SCL co-ordination language [3]. They sought to integrate imperative code chunks within a structured parallel framework. As an example, the P^{3}L language

core includes programming paradigms like pipelines, task farms, iterative and data parallel skeletons. Skeletons in P^3L can be used as constructs of an explicitly parallel programming language, actually *as the only way to express parallel computations.*

Later on, all experiences assessed with P^3L have met into the SkIE language and its compiler [4]. In SkIE existing sequential codes can be used to instance skeletons with little or no amendment to the sources; it supports several guest sequential and parallel languages (C, C++, Fortran, Java, HPF) within the same application. A SkIE program is basically a composition of skeletons. Also, they are equipped with a compositional functional semantics. They behave like higher-order functions which can be evaluated efficiently in parallel. Furthermore, the skeletons functional and parallel semantics enabled the optimization of programs by means of performance-driven source-to-source code transformations [5,6,7].

As SkIE concerns, several real world applications* have been used as test-bed to validate the effectiveness of the programming environment [8]. A lack of expressivity emerged for some of them. In principle, the skeletal approach is not particularly targeted towards a class of applications. However, we experienced that some applications can be straightforwardly formulated in terms of skeleton composition, others needs a greater design effort. The boundary between the two classes depends on many factors, among the others, the particular programming environment and the skeleton set chosen for applications development. Anyway, some common flaws may be recognized in both SkIE and other research group works (see also [9]):

(i) The selection of skeletons to make available in the language skeleton set is quite critical design issue. Despite several endeavors to classify and close the parallel programming skeleton set [10], in many cases during application development we experienced the need of the "missing skeleton", or at least the missing functionality for an existing skeleton.

(ii) Many parallel applications are not obviously expressible as instances of (nested) skeletons, whether existing or imagined. Some have phases which require the use of less structured or ad-hoc interaction primitives.

(iii) Although all kind of languages may be equipped with a skeletal super-structure, skeletal languages has been historically designed in a functional programming style fashion [4,11]. In this setting non functional code is embodied into the skeletal framework by providing the language with wrappers acting as pure functions. Actually, the fully functional view (by its very nature) does not enhance programmer control over data storage that is a feature that may happen to be useful in the design of applications managing large, distributed, randomly accessed data sets.

Indeed, the role of skeletons in the programming language is evolved and matured along past decade. Such evolution yearning to defeat (among the others) described lacks in skeletal languages expressiveness while preserving their ease of use. In particular skeletons loose their "exclusiveness" on parallelism exploitation. The new skeletons' role has led to the exploration of several scenarios:

*In the area of massive data mining, computational chemistry, remote sensing and image analysis, visual and numerical computing.

Skeletons as Design Patterns. A design pattern *per se* is not a programming construct, as happened for the skeletons. Rather, it can be viewed as "recipe" that can be used to achieve different solutions to common programming problems. The parallel skeleton support may be implemented using a layered, OO design [12]. Parallel skeletons can be declared as members of proper skeletons/patterns. Exploiting standard OO visibility mechanisms, part of the framework may be made visible to the programmer in such a way he can perform different tasks: fine performance tuning, introduction of new, more efficient implementation schemes, etc. [13,14].

Skeletons as Extension. Skeletons may be used to extend existing programming languages or programming frameworks (e.g. C + MPI) that are already able to exploit parallelism. Several recent programming frameworks may be numbered among this category, among the others: `SKElib` [15] that extends C language with SkIE-like skeletons and enable the programmer to use standard Unix communication mechanisms. `Skil` [16] extend C++ language providing the programmer a SPMD environment with task (pipeline and farm) and data parallel (map, fold, ...) skeletons, that are seen as collective operations. `Lithium` [17] is the first *pure Java* structured parallel programming environment based on skeletons and exploiting macro data flow implementation techniques [18]. `eSkel` [9] is a library which adds skeletal programming features to the C/MPI parallel programming framework. It is a library of C functions and type definitions which extend the standard C binding to MPI with skeletal operations.

eskimo has been influenced by both previous approaches. eskimo extends the C language with "proto-skeletons" or constructs, which represent skeletons' building blocks. Skeletons does not really exist in eskimo program as language elements, rather they are particular programming idioms.

3. eskimo: A New Skeletal Language

eskimo is a parallel extension of C language based on shared address programming model. The target architectures for the language are Beowulf class machines, i.e. POSIX boxes equipped with TCP/IP networks. In this setting, eskimo is conceived to be a framework to experiment the feasibility of the skeletal approach with dynamic data structures in parallel programming.

The basic idea behind eskimo is that a programmer should concentrate on co-designing his data structures and his algorithms. Moreover, in order to obtain a high-performance application, the programmer would structure its application properly, and eventually suggest to run-time important information about algorithm data access patterns. eskimo run-time takes care of all other details like process scheduling and load balancing. eskimo run-time support is based on a software distributed shared memory. Notably it is not yet another DSM, rather it relies on DSM already known technologies to experiments the co-design of dynamic data structures and parallel programming patterns enforcing locality in the distributed memory access. We outline the main features of eskimo as follows:

Abstraction. eskimo is a skeleton based programming language. It aims to simplify programming by raising the level of abstraction, providing to the programmer performance and portability for its applications. In order to convey this simplicity to programmers we must be careful not to bundle it with an excessive conceptual

baggage. At this end we enriched C language in such a way the language extension fairly raises the level of abstraction. The main sources of abstraction regard *data structures, the flow of control, and the interaction between them.* All abstractions rely on solid concepts like concurrency and abstract data types.

Expressiveness. We propose a structured programming environment that allows the programmer to deal with (dynamic) spread shared data structures. In particular, the programmer deals with an abstraction of data structures represented as a single entity (as in [19]). These parts are kept consistent by the run-time support following a (very) lazy memory consistency model (see sec. *3.3*). The chosen consistency model enforces the high-level approach of the language since it enables to read/write data objects avoiding the need of explicit low-level synchronization primitives (like locks and barriers). In this setting the skeleton is no longer a ready-made object of the language (e.g. an high-order function), rather it is a code pattern build directly by the programmer using language primitives. Ad-hoc parallel pattern may be coded using both eskimo and other libraries primitives.

Framework and design principles. eskimo main target architecture class are Beowulf class clusters. Such architectures, that are becoming pretty popular due to their limited cost, presents several difficulties in drawing good steady performance from applications (particularly dynamic ones). Following the nature of target architecture class, eskimo exposes to the programmer a (virtual) shared NUMA address. The programmer is required to make decisions about the relationship among data structures (e.g. locality) but not to deal with all cumbersome facets of data mapping. The underlying design principle consists in considering preferable a programming environment on which performance improves gradually with increased programming effort (taking advantage from a deep application knowledge) with respect to one that is capable of ultimately delivering better performances but that requires an inordinate programming effort. This can consist in either programming each detail of the application (as in low-level approaches) or expressing the application attempting to use a fixed set of ready-made parallel paradigms.

In summary eskimo extends the C language with three classes primitives: flows of control management (sec. *3.1/3.2*); *Shared Data Types* declaration, allocation and management (sec. *3.3*); shared variables management (sec. *3.4*).

3.1. Exploiting Parallelism in eskimo

The parallelism is exploited through concurrency. The minimal unit for concurrency exploitation is the C function. Just as in a serial program, an eskimo program starts as a single control flow. In any part of the program, the programmer may split the flow of control through the asynchronous call of a number of functions; such flows must, sooner or later, converge to a single flow of control. The basic primitives managing program flow of control behave like Dennis' fork/join, we call them *e-call/e-join.* Also, we call *e-flows* eskimo flows of control. *e-flows* share the virtual memory address space. The relationship among *e-calls, e-joins* and *e-flows* is discussed in section *3.2* and intuitively sketched in Fig. 1 a).

e-call/e-join primitives enable the programmer to set up a dynamic and variable number of *e-flows*, that is a pretty important feature dealing with dynamic data structures (in particular linked data structures as lists and trees). Almost all

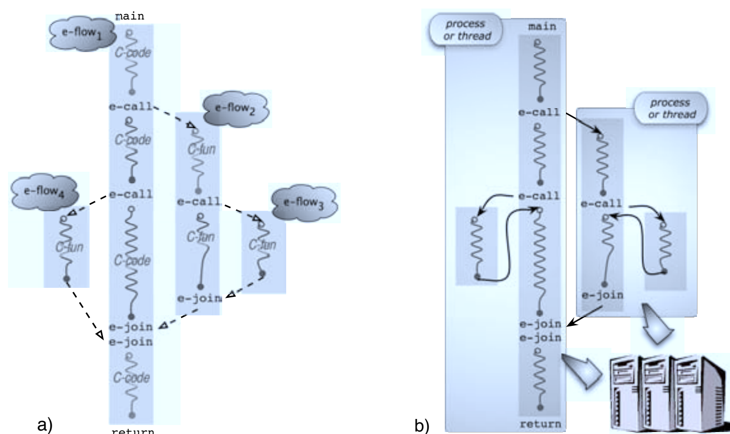

Fig. 1. An eskimo program execution intuitive view. a) Relationship among *e-calls, e-joins* and *e-flows* (grey boxes). b) A possible execution of the program.

interesting algorithms on these data structures explore them in a recursive fashion, following the Divide&Conquer ($\mathcal{D}\&\mathcal{C}$) paradigm exploiting a variable concurrency availability along the program run lifespan.

Actually $\mathcal{D}\&\mathcal{C}$ has been already present in classical skeleton sets. This skeleton, as others (e.g. scan and map), has been often interpreted as a collective operation on a given data structure by the skeleton community [6,21]. In the best case this originated a family of variants for each skeleton, each of them optimizing a particular behavior of the algorithm on the data structure [22]. In other cases, application programmers have been compelled to match a given shape of the data structure with a given behavior of the skeleton (that is likely to be a frustrating task). eskimo approaches the problem from a lower-level viewpoint: it enables the co-design of algorithms and data structures as in the very classical sequential programming. There is no ready-made $\mathcal{D}\&\mathcal{C}$ skeleton in eskimo, rather there are all ingredients to build it in such a way the programmer may express the suitable variant of the skeleton for its data structure. Then, the language run-time tries to understand from the ingredients and from the programmer hints what is the expected parallel behavior for the particular variant, and even in case hints are wrong or missing the program does not loose its correctness (even if we cannot expect from it an optimal performance).

3.2. Concurrency and Flows of Control

Actually *e-flows* do not necessarily match any concrete entity at eskimo run-time support level or its underneath run-time layers (as for example threads or processes). In particular an *e-call* has to be considered as the declaration of a "concurrency capability" with respect to a given function instance: an *e-called* function instance might be both concurrently executed or sequentialized with respect to the caller function. In the former case, the matching *e-join* represents the point along the execution unfolding where a called *e-flow* must converge in to the caller *e-flow*.

The two *e-flows* coming out from an *e-call* may be executed in parallel, interleaved or serialized in any order depending on the algorithm, the input data and the system status. Actually, as *e-call/e-join* primitives denote in the algorithms points where it is possible to proceed in more than one way, even for the same input data set, they are *non-deterministic* primitives. As an example, eskimo program sketched in Fig. 1 a) exploits four *e-flows* that represents its top concurrency degree. However, as shown in Fig. 1 b) the language run-time might choose to halve the concurrency degree and project the four *e-flows* in two really concurrent entities. Serialized *e-flows* do not lead to any overhead related to parallelism exploitation (thread creation, activation, synchronization).

As might be expected, *e-flows* have to be mapped on concrete concurrent entities (i.e. processes or threads) at the language run-time level. This mapping is sketched on-the-fly, step by step (in a distributed fashion) by eskimo run-time; in correspondence of an *e-call* the calling *e-flow* map the new *e-flow*. Actually, the mapping process must answer to two key questions:

1 Does the new *e-flow* be really concurrently executed (either in parallel or interleaved) or it has to be serialized with respect to the calling one?
2 In case of concurrent execution, does the fresh *e-flow* be locally or remotely spawned?

In both cases the language run-time looks for a tradeoff between opposite needs. In the former case, it tries to maintain the amount of concurrent flows in the system within acceptable bounds: enough to exploit the potential speedup of the system, but not to much in order to avoid unnecessary overheads (in time and memory space) due to parallelism management. In the latter case, the run-time tries to balance workload on PEs while keeping data locality as much as possible.

We highlight that the relationship among *e-flows* is slightly more abstract than concurrency. Actually different *e-flows* encompasses parts of an application which might be concurrent each other, but in some cases they are not. Consider some *e-flows* called in a sequence that establish a (direct or indirect) precedence relation of one over another; or some *e-flows* sequentialized in to the same concurrent entity. In both cases *e-flows* do not correspond to actual concurrent code.

eskimo *e-call/e-join* are the basic primitives to create and destroy an *e-flow*. In addition to them, eskimo provides the programmer with their generalization, i.e. *e-foreach/e-joinall*. They work basically in the same way but, as shown in Fig. 2 b), they can create and destroy an arbitrary number of *e-flows*. Since *e-flows* created by means of *e-foreach* have no data dependencies one each other, they can be non-deterministically executed. Complementary *e-joinall* non-deterministically waits the completion of all *e-flows* created by the matching *e-foreach*. *e-foreach/e-joinall* have an additional freedom degree with respect to a sequence of *e-call/e-join*: the order in which *e-flows* are mapped/scheduled and joined. From the run-time viewpoint, this turns in some additional advantages with respect to the basic case. In particular, at *e-foreach* time the language run-time knows how many and which *e-flows* have to be executed, and may choose their execution order depending on data availability. eskimo run-time uses this possibility to enhance locality of data accesses by running first tasks that are likely to have needed data (or a part of

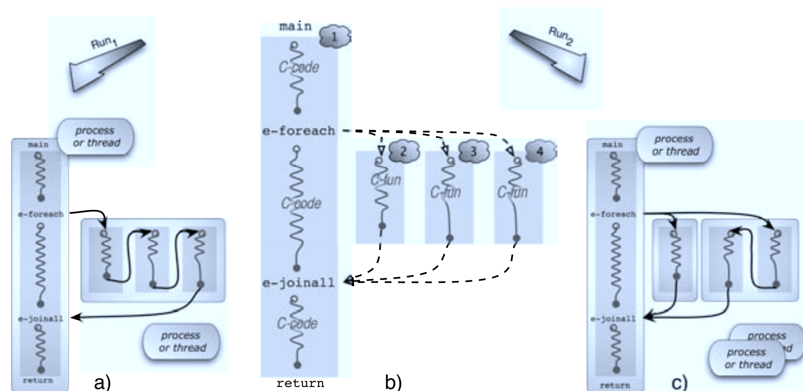

Fig. 2. An eskimo program resulting in two different mapping and scheduling in different runs.

them) already present (or cached) in the PE. In the same way, the run-time joins *e-flows* in the order they complete instead than the program order. Therefore, the same program may be subjected to different mapping and scheduling decisions on different runs (even on the same input data). As an example, in Fig. 2 a) and c) are sketched two runs of the same program that have been subjected to different mapping and scheduling decisions.

3.3. Sharing Memory among Flows of Control

eskimo language is specifically thought for NUMA parallel computing frameworks, in particular for distributed memory architectures. These architectures naturally supply a very efficient memory access to local memory and a more expensive access to remote memory. The language abstracts these memory categories, but it does not obscure the differences between them. eskimo language exposes to the programmer two memory spaces: *private* and *shared*.

Private and shared variables are distinguished by their types. Shared variables must have a *Shared Data Types* (SDT). SDTs are obtained by instancing a fixed set of *Shared Abstract Data Types* (SADT), i.e. parametric types, including arrays, k-trees and regions.[†] All variables with a different type are private. However, not all C variables are allowed as private variables, in particular C global variables are (with some exception) forbidden into eskimo programs. These must be substituted with shared variables. The relationship between *e-flows* and memory classes is the following:

- Each *e-flow* has its private address space; thus private variables are always bound to the *e-flow* where they have been declared. Private variables be accessed without any eskimo mediation.
- All *e-flows* can access the shared address space, that is spread across the system. Shared variables may be accessed by several *e-flows*, thus might be

[†] A contiguous region of memory that is shared but not spread.

subjected to concurrent accesses. These accesses are regulated by the eskimo run-time.

Overall, SADTs and their constructors provide a device to make sharable any C data type by building SDTs. These are data structure containers allowing the access to contained data (i.e. private or shared variables) by several *e-flows*. In case of trees and arrays the container is structured and holds many contained data in a lattice; such data is distributed across the virtual architecture thus it may have a total size greater than memories of the single PE. Shared variables (or part of them) may be referenced across *e-flows* by using *references*, i.e. void pointers into shared address space. References, are the basic mechanism to pass shared variables across *e-calls*. Pragmatically, eskimo references and shared variables would provide globally addressable data structures matching the same relation between C pointers and variables.

All SDTs are designed to be concurrently accessed by *e-flows*. SDTs may be statically or dynamically (and incrementally) allocated, in particular $k-$tree nodes must dynamically allocated by means of a language primitive (`e_node_add`). Beyond trees, the programmer may build any shared linked data structure by using references. Shared variables obey to DAG consistency [20] and can be accessed through eskimo primitives (see next section). That basically means two different *e-flows* may have a non coherent view of a given address in the shared memory space; the coherence is then reconciled at *e-join* time. Pragmatically it means different *e-flows* must write different shared address locations.

3.4. Reading and Writing Shared Variables

Most software DSMs rely on MMU and operating system traps in order to make transparent the access of remote data. However, the interrupt cost, associated with receiving a message has been proved to be the largest component of the slow remote latency, not the actual wire delay in the network or the software implementing the protocol [23]. We followed a different approach, shared variables, differently from private ones, can not be read and written directly. In order to access to a shared variable the programmer must explicitly bind it to a private pointer (see Fig. 3 line 4). The technique avoid the use of operating system traps. In addition it has a pragmatic importance: since accessing a shared variable may be expensive the programmer is required to evaluate really needed accesses[‡] eskimo provides a pair of language primitives enabling the access to a shared variable: `r` (for read-only) and `rw` (for read/write). Both `r` and `rw` take as argument a reference and return a void pointer that can be used (after a suitable type cast) to access the shared variable value.

4. Skeletons and their expected pay-back

eskimo provides the programmer with a family of constructs that specialize *e-foreach*. Each of them run through elements of a given type of shared variables. Currently there are three types of *e-foreach* constructs:

[‡] Anyway, transparency may be achieved by preprocessing/instrumenting the user code.

`e_foreach_child` run through non-null children of a spread tree node;
`e_foreach_cell` run through cells of a spread array;
`e_foreach_refset` run through a given set of references to shared variables or elements of them (as for examples the set of tree leafs, or array cells).

The *e-foreach* constructs introduce a form of data parallelism, which is based on on domain decomposition principle. eskimo allows the programmer to dynamically define and decompose the domain. In fact, the parallel application of a function to (sub)set of children in a tree may be interpreted as a form of data parallelism (where the dynamic domain of children is decomposed). As result, also $\mathcal{D}\&\mathcal{C}$ paradigm may be interpreted as a special (dynamic, recursive) case of data parallelism. As an example consider the Barnes-Hut application code fragment shown in Fig. 3.

eskimo offers to the programmer the possibility to store data using a flat (arrays) or hierarchical (trees) structure, then offers two standard parallelization paradigms for the two cases: forall (or map) and $\mathcal{D}\&\mathcal{C}$. The two paradigms may be freely interleaved and are both introduced at the language level by just one primitive (family), namely the *e-foreach*. In case the application is not obviously expressible as instances of proposed paradigms or their interleaving, the programmer may ad-hoc parallelize it by using the standard C language enriched with *e-calls/e-joins*.

The basic idea under eskimo is to abstract both (dynamic) data structures and application flow of control in such a way they result *orthogonalized*. Creating an *e-flow* either an *e-called* function may be migrated to the PE holding the data it need or vice-versa. Clearly, the principal decision is to evaluate which is the better choice in each case. In such task the run-time takes in account a number of elements: the system status (load balancing), the shared memory space status and the programmer hints. Several policies may be implemented using such informations. Since Beowulf class cluster exploits an unbalanced communication/computation power tradeoff, the run-time tries to schedule an *e-flow* on the same PE of the data it needed and takes in account load-balancing only as secondary constraint. Clearly, the key issue here is to know in advance what data a function will access. The run-time uses two kind of informations: (i) The data blocking. The run-time makes scheduling decisions only when a function parameters cross the current data segment boundary. As result *e-flows* accessing to data in segment are sequentialized enhancing locality and coarsening computation grain. (ii) The programmer hints. We assumed that the first parameter of each *e-called* function (that must be a reference) is considered as a dominant factor, the run-time expect the majority of data allocated "close" to data referenced by it. A wrong/missing information has no impact on program correctness but a great impact on its performance.

5. Implementation and Experiments

A prototypal implementation of eskimo already runs on Linux clusters. It has a layered implementation design: each layer provides mechanisms and policies to solve some of parallel programming support issues. In particular *e-flows* mapping and scheduling as well as shared data structure mapping is implemented in the top tier. At this level many informations about the system status may be accessed (e.g. PEs load, memory load, etc.). These information are maintained by a lower level

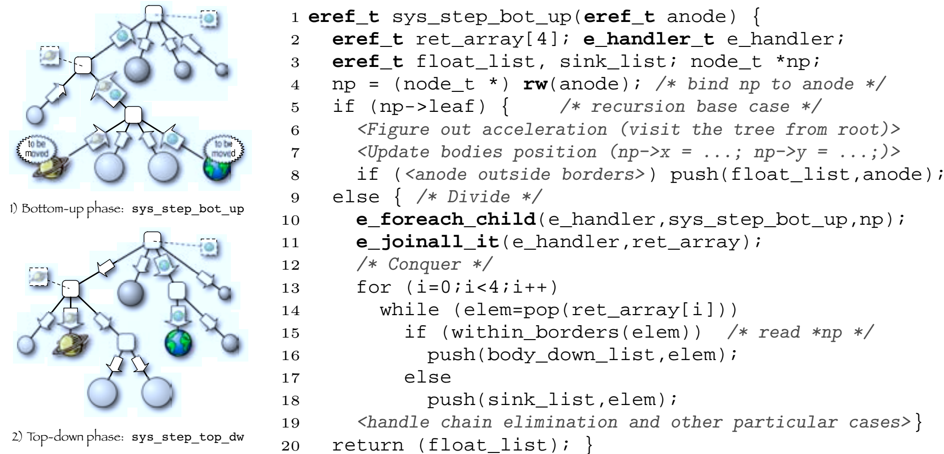

Fig. 3. A n-body system step in two phases and the eskimo pseudocode of the first phase.

tier that basically wrap the communication stack (currently the TCP stack). The run-time exploits both multiprocessing (inter-PEs) and multithreading (intra-PEs).

We tested eskimo against Barnes-Hut application. It is a good example to study because present non-trivial performance challenges due to irregular and dynamically changing nature of data structures, and it is relatively small and manageable. We developed three different versions of the Barnes-Hut algorithm: 1) The sequential version (seq), implemented in C language. It is a reduction to the bare bones of the original Barnes-Hut code in the bidimensional space. 2) The eskimo version. It is obtained from the sequential version substituting recursive calls with recursive *e-foreach*es. The n-body system evolves in two phases: (references to) bodies leaving their current quadrant are first lifted to the smallest quadrant including both source and target positions, then they are pulled down to target quadrant (see Fig. 3 left). The main data structure is a spread quad-tree. 3) The C+MPI version. The body data is partitioned among nodes. The hierarchical relationship among bodies are maintained in a forest of trees. Each tree of the forest is linked to a "root" tree replicated in each PE. The structure simulates a spread tree with the top part cached in each PE for a faster access. A PE maintains coherent the top part of the tree and read/write other parts of the tree by exchanging messages with other PEs.

We tested the three application versions on two different datasets: *plummer* and *uniform*. The two distribution may be hierarchically represented by a strongly-unbalanced and fairly-balanced trees respectively. Table 1 shows the performance and speedup figures of the three versions of the algorithm on the SMP cluster for four different datasets. The eskimo version of the code results as fast as the MPI version for the uniform dataset and slightly but significantly better for the plummer dataset (highlighted in the table). The point here is that the performance MPI version does not scale up with the number of PEs for the plummer dataset. The tree unbalanced leads to a heavy load imbalance in the MPI version. Our version of the MPI code does not include a dynamic load balancing strategy but

dataset·#bodies	Pl·10k	Pl·20k	Un·10k	Un·20k	Pl·10k	Pl·20k	Un·10k	Un·20k	Optimal
	time				speedup				
seq	6.47	14.53	4.80	2.34	–	–	–	–	–
MPI 1x2-way SMP	6.60	14.36	2.54	1.27	0.9	1.0	1.9	1.8	2
MPI 2x2-way SMP	6.64	14.58	1.50	0.75	0.9	1.0	3.2	3.1	4
eskimo 1x2-way SMP	5.30	13.20	2.45	1.33	1.2	1.1	1.9	1.8	2
eskimo 2x2-way SMP	4.10	8.10	1.55	0.77	1.6	1.8	3.1	3.0	4

Table 1. Barnes-Hut performances (secs) and speedups on several Plummer (Pl) and uniform (Un) datasets on a SMP cluster (2-way 550MHz Pentium III).

fixes the data distribution during the first iteration. It is certainly possible to add a dynamic balancing strategy in the MPI code (even if not so easy), but it has to be explicitly programmed by the application programmer and specifically tailored for the problem. eskimo code instead can be written without any concern for load balancing and data mapping. As matter of fact the sequential version is just 300 lines of code, the eskimo version 500 and the MPI version 850.

6. Conclusions, Related Work and Acknowledgements

eskimo programming framework is designed and developed from scratch but it has several analogies with a number of well known research works. Among the others it is worth to mention *Cilk* [24] and *Athapascan* [25]. In particular, eskimo inherits memory consistency model from *Cilk*. The layered implementation design and the idea to make customizable the scheduling is quite similar to *Athapascan*.

Nevertheless, eskimo differs from them in a number of design issues. It does not use work stealing, that has load-balancing as main target; it tries to exploit the combined mobility of data and computations with the aim of reducing network traffic by using programmer insight on frequent data access patterns. It implements dynamically allocable dynamic data structures (trees). These are designed to be blocked in order to reach an acceptable working grain for the target architecture class in a transparent way. It does not rely on any preprocessing; it is simply a C library. It is conceived for loosely coupled architectures. It does not rely on any shared stack (e.g. cactus task) for function calls: the stack never moves across PEs and may be optimized by the standard C compiler (e.g. in-lining). Last but not least, it abstract flow of control and data management in a skeletal perspective.

eskimo exploits a good performance on a non-trivial problem using dynamic data structures. However, it is a preliminary work and a number of features remain to be fixed. In particular, we believe that the adoption of an OO (as in [19,17]) host language would greatly simplify the SADT definition and implementation as well as the language syntax.

This work has been partially supported by the Italian Space Agency (ASI-PQE2000 Project) and by the National Research Council (Agenzia 2000 Project).

References

[1] M. Cole. *Algorithmic Skeletons: Structured Management of Parallel Computations*. Research Monographs in Parallel and Distributed Computing. Pitman, 1989.

[2] B. Bacci, M. Danelutto, S. Orlando, S. Pelagatti, and M. Vanneschi. P^{3}L: A Structured

High level programming language and its structured support. *Concurrency Practice and Experience*, 7(3), May 1995.
[3] J. Darlington, Y. Guo, Y. Jing, and H. W. To. Skeletons for structured parallel composition. In *Proc. of 15th PPoPP*, 1995.
[4] B. Bacci, M. Danelutto, S. Pelagatti, and M. Vanneschi. SkIE: a heterogeneous environment for HPC applications. *Parallel Computing*, 25(13–14), Dec. 1999.
[5] M. Aldinucci and M. Danelutto. Stream parallel skeleton optimization. In *Proc. of 11th PDCS*. IASTED/ACTA press, Nov. 1999.
[6] M. Aldinucci, S. Gorlatch, C. Lengauer, and S. Pelagatti. Towards parallel programming by transformation: The FAN skeleton framework. *Parallel Algorithms and Applications*, 16(2–3), 2001.
[7] M. Aldinucci. Automatic program transformation: The Meta tool for skeleton-based languages. In *Constructive Methods for Parallel Programming*, Advances in Computation: Theory and Practice, chapter 5. Nova Science Publishers, 2002.
[8] G. R. Joubert, A. Murli, F. J. Peters, and M. Vanneschi, ed. *Parallel Computing: Advances and Current Issues. Proc. of ParCo2001*. Imperial College Press, 2002.
[9] M. Cole. Bringing skeletons out of the closet. Technical report, Uni. Edinburgh, 2002. (`http://www.dcs.ed.ac.uk/home/mic/eSkel/eSkelmanifesto.ps`).
[10] S. Pelagatti. *Structured Development of Parallel Programs*. Taylor&Francis, 1998.
[11] J. Sérot and D. Ginhac. Skeletons for parallel image processing: an overview of the SKIPPER project. *Parallel Computing*, 28(12), Dec. 2002.
[12] S. MacDonald, J. Anvik, S. Bromling, J. Schaeffer, D. Szafron, and K. Tan. From patterns to frameworks to parallel programs. *Parallel Computing*, 28(12), Dec. 2002.
[13] M. Danelutto. On skeletons and design patterns. In *Parallel Computing: Advances and Current Issues. Proc. of ParCo 2001*. Imperial College Press, 2002.
[14] F. A. Rabhi and S. Gorlatch, ed. *Patterns and Skeletons for Parallel and Distributed Computing*. Springer-Verlag, 2002.
[15] M. Danelutto and M. Stigliani. SKElib: parallel programming with skeletons in C. In *Proc. of Euro-Par 2000*, LNCS n. 1900, Sep. 2000.
[16] G. H. Botorog and H. Kuchen. Skil: An imperative language with algorithmic skeletons for efficient distributed programming. In *Proc. of 5th HPDC*. IEEE, 1996.
[17] M. Aldinucci, M. Danelutto, and P. Teti. An advanced environment supporting structured parallel programming in Java. *Future Generation Computer Systems*, 19(5), 2003.
[18] M. Danelutto. Efficient support for skeletons on workstation clusters. *Parallel Processing Letters*, 11(1), Mar. 2001.
[19] H. Kuchen. A skeleton library. In *Proc. of Euro-Par 2002*, LNCS n. 2400, 2002.
[20] R. D. Blumofe, M. Frigo, C. F. Joerg, C. E. Leiserson, and K. H. Randall. An analysis of dag-consistent distributed shared-memory algorithms. In *Proc. of 8th SPAA*. ACM, Jun. 1996.
[21] S. Gorlatch. Send-Recv considered harmful? Myths and truths about parallel programming. In *Proc. of PACT 2001*, LNCS n. 2127, 2001.
[22] C. A. Herrmann. *The Skeleton-Based Parallelization of Divide-and-Conquer Recursion*. PhD thesis, FMI Uni. Passau, Germany, 2001.
[23] Z. Radović and E. Hagersten. Removing the overhead from software-based shared memory. In *Proc. of Supercomputing 2001*. ACM, Nov. 2001.
[24] R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C. E. Leiserson, K. H. Randall, and Y. Zhou. Cilk: An efficient multithreaded runtime system. In *Proc. of 5th PPoPP*, 1995.
[25] F. Gallilée, J.-L. Roch, G. G. H. Cavalheiro, and M. Doreille. Athapascan-1: On-line building data flow graph in a parallel language. In *Proc. of PACT '98*. IEEE, 1998.