, , 1–27 ()
© Kluwer Academic Publishers, Boston. Manufactured in The Netherlands.

On Migrating Threads *

BERND MATHISKE mathiske@fb.sony.de
FLORIAN MATTHES matthes@informatik.uni-hamburg.de
JOACHIM W. SCHMIDT j_schmidt@informatik.uni-hamburg.de
Fachbereich Informatik, Universität Hamburg,
Vogt-Kölln Straße 30, D-22527 Hamburg, Germany

Editor: Amihai Motro

Abstract. Based on the notion of persistent threads in Tycoon [30], we investigate thread migration as a programming construct for building activity-oriented distributed applications. We first show how a straight-forward extension of a higher-order persistent language can be used to define activities that span multiple (semi-) autonomous nodes in heterogeneous networks. In particular, we discuss the intricate binding issues that arise in systems where threads are first-class language citizens that may access local and remote, mobile and immobile resources.

We also describe how our system model can be understood as a promising generalization of the more static architecture of first-order and higher-order distributed object systems. Finally, we give some insight into the implementation of persistent and migrating threads and we explain how to represent bindings to ubiquitous resources present at each node visited by a migrating thread on the network to avoid excessive communication or storage costs.

Keywords: distribution threads migration agents workflows persistence higher-order languages bindings dynamic linking

1. Introduction and Rationale

Most of the work on migrating threads described in this article has been carried out in the context of the European ESPRIT basic research project FIDE (fully integrated data environments). The rationale behind the FIDE project is to improve significantly the productivity in the process of building integrated, data-intensive applications. The overall goal of FIDE is to develop a consistent and small set of orthogonal language and system concepts to eliminate the historical mismatches and overlaps between independently developed component technologies.

While the initial contribution of FIDE has been the development of persistence and type system technology to overcome the mismatches between programming language and database technology, the work described here addresses the overlaps and mismatches between distributed programming, transaction management and workflow management.

* This work was supported in part by the Commission of the European Communities under grant number ISC-CAN-080 CIS, *Activity Modelling and Object Technology for Cooperative Information Systems.*

More specifically, we propose to generalize the well-understood programming concept of threads [38] to also cover persistent and migrating threads which can be used as primitive building blocks to implement long-term and distributed cooperative activities. In a recent paper [30], we describe the notion of persistent threads while this article reports on our work to migrate such persistent threads between heterogeneous nodes in local or wide area networks.

The analogy with the concept of remote procedure calls may help to clarify the rationale behind our work. Remote procedure calls and migrating threads both demonstrate how to reduce successfully the complexity of distributed systems by a generalization of a well-established non-distributed programming concept:

- Programmers do not need to learn a new programming abstraction.
- Local software architectures and libraries can be scaled easily to distributed software architectures.
- Tools can be provided to optimize the mapping from high-level language abstractions to low-level communication mechanisms.
- The generalized programming abstraction interacts well with other programming language concepts like static typing, modularization and parameterization.

Our work on migrating threads reveals a strong synergy between concepts, language constructs and technologies which support persistence through time with their counterparts to support distribution through space. This synergy will become obvious in the rest of this article which is organized as follows.

First, we present migrating persistent threads as a generalization of conventional thread concepts and motivate them as a valuable programming abstraction for the implementation of long-term distributed activities. We then introduce a terminology for data, code and thread bindings (Section 3) and basic thread operations as well as related typing aspects (Section 4). This is the basis for a comparison of alternative models for distributed programming in Section 5. These range from a client-server oriented programming style via higher-order languages and object migration approaches to migrating threads. In Section 6 we discuss the difficult binding issues arising in the context of thread migration. Based on our classification of thread resources we describe how to achieve adequate binding support for standard programming situations. In Section 7 we give some insight in our implementation of migrating threads available uniformly on Unix, Windows and MacOS platforms.

2. Thread Persistence and Mobility

A thread describes a single sequential flow of control in a program. Having multiple threads in a program means that at any instant the program has multiple points

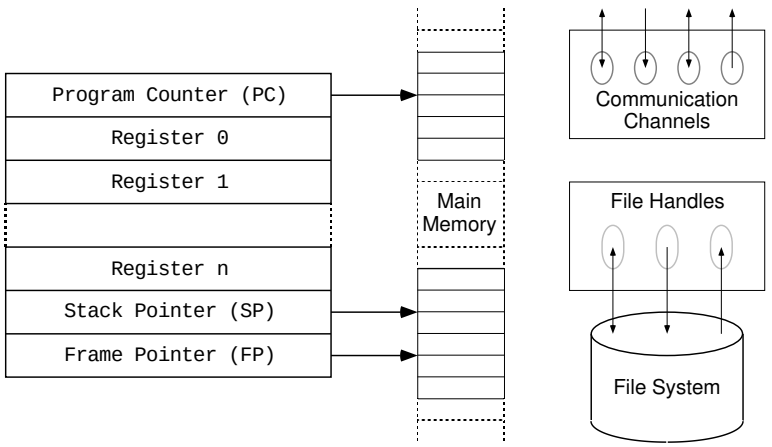

Figure 1. Evaluation state and resources of a local volatile thread

Table 1. Thread-based implementations of cooperative activities

Cooperation Mode	same time (single session)	different time (multiple sessions)
same location (single address space)	local volatile threads	local persistent threads
different location (multiple address spaces)	migrating threads	migrating persistent threads

of execution, one in each of its threads. Unlike operating system processes, multiple threads execute within a common address space, permitting multiple threads to access shared resources (see Figure 1). Threads are also known as *lightweight processes* [42], [34], [19] as they require significantly fewer resources (mostly time and storage space) than operating system processes. This concerns in particular creation and context switches.

Today, threads are a well-established and standardized [38] programming abstraction suitable for short-lived and non-distributed applications. For example, loosely coupled high-level activities like text editing, spell checking, typesetting, and screen refreshing can be mapped to separate concurrent threads that execute different code against shared data. Moreover, there is a rich repertoire of explicit and implicit synchronization mechanisms for thread coordination and communication (semaphores, mutexes, condition variables, channels, message queues, rendezvous', transactions).

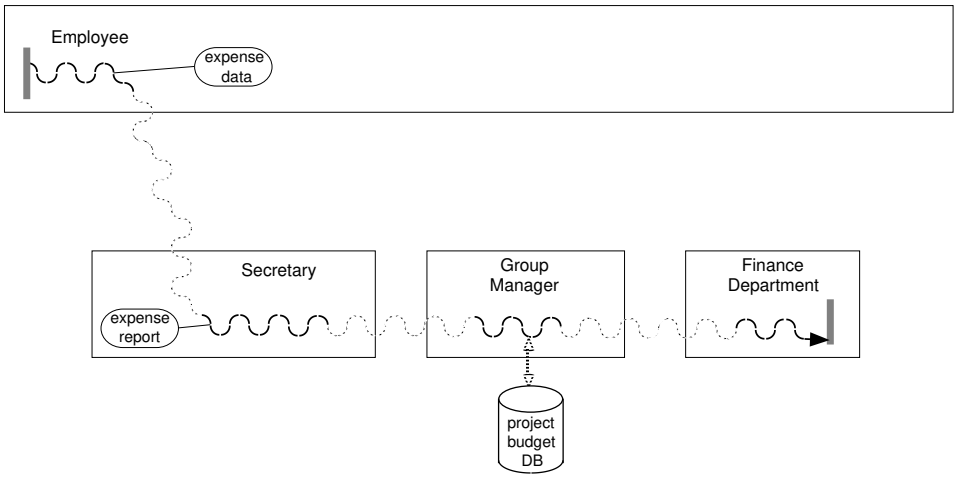

Figure 2. Implementing a distributed activity by thread migration

However, the traditional thread concept exhibits severe limitations for cooperative work as required by data-intensive applications, as it covers only one segment of the space vs. time matrix depicted in Figure 1:

- The lifetime of a thread is limited by its enclosing operating system process (i.e., a thread is a *volatile* and not a persistent value).
- The location of a thread is limited by a single address space.

As described in [30], a generalization towards *persistent threads* makes it possible to map also long-term activities (business processes or workflows) to threads. For example, the handling of a customer request by a clerk in a service department can be modeled by a thread that runs for several days or even weeks, thereby outlasting multiple activations of the same program. In a type-complete persistent language, persistent threads can be stored, for example, as attributes of database tables and they can be manipulated by user-defined queries.

If one adopts an activity-oriented view of distributed applications, it is desirable to be able to express distributed workflows directly by *migrating threads* that span multiple network nodes and independent databases. For example, Figure 2 and the associated Tycoon script in Figure 3 describe an expense report activity in a company by means of a single thread. This thread migrates through different departments (sites) and carries around information about its past history. Being invoked by an employee, it is equipped with various bindings to mobile resources like the attributes of the expense report and immobile resources like a project budget

databases maintained by the group manager. Passing through the secretary's and the group manager's site, it eventually reaches the finance department where a money transfer is initiated.

In a conventional system, activity persistence is implemented by writing the full thread state to a file or a database. To continue an activity after a process restart, its stored state has to be reloaded and a new thread has to be created in order to resume code execution at the last instruction executed in the previous session. Similar techniques have to be applied for activity migration. This "manual" approach to thread persistence and migration fails for non-trivial long-term and distributed activities due to the following technical difficulties:

- Bindings between variable names in the code (e.g., *expenseReport, projectBudgetDB* in Figure 3) and their associated entities have to be reestablished. Such thread bindings are discussed in more detail in Section 3.

- Depending on the power of the language at hand (loops, function calls, recursion, exception handling, etc.) it may be very difficult to recreate a thread that is in the correct execution state (program counter, evaluation stack, exception handler).

- It is often necessary to transmit the state and code information of a thread separately and to apply ad-hoc dynamic code binding mechanisms at the receiver side prior to thread restart.

In our Tycoon system and language, persistence and mobility is the default for all data, code and thread entities. Only immobile and volatile resources require special treatment by the programmer (see Section 6). This relieves the programmer from low-level thread implementation details and turns threads into a suitable programming primitive for distributed, data-intensive applications.

3. Thread Bindings

In this section we introduce a terminology for the description of data, code and thread bindings (c.f. [30]) that we use in the rest of this article to describe the Tycoon thread semantics, to discuss selected thread implementation aspects and to compare different models for distributed programming.

A binding is an association between a name and a computational entity from a specific semantic domain [41], [32]. We also say that a name is bound to a computational entity. An environment is a (possibly ordered) collection of bindings. Names are used to identify entities in an environment. Different names can be bound to the same entity (sharing, aliasing).

Entities can be atomic (like integers or booleans) or structured (like records, objects or functions). Structured entities typically consist of environments. For example, the fields of a record lead to bindings from field names to other entities. Therefore, bindings can be used to model (recursive) relationships between entities.

```
migrate to Employee do
  repeat
    let data = getExpenseDataFromUser()
    migrate to Secretary do
      let expenseReport=compileReport(data)
    end
  until valid(expenseReport)
end
migrate to GroupManager with remote
   projectBudgetDB :ProjectBudgetDB
do
   update(projectBudgetDB, expenseReport.total)
end
migrate to FinanceDepartment do
   transferMoney(expenseReport.total)
end
```

Figure 3. Tycoon script for a migrating thread (expense report)

Entities can be flat (like records) or nested (like functions in Algol-like languages). In a nested entity, names bound in a global outer environment are automatically visible in a local inner environment.

The binding of names to structured entities naturally leads to the concept of a transitive closure of bindings that underlies many persistence and migration models: any entity reachable through a chain of bindings from a persistent (mobile) entity becomes persistent (mobile), too. This approach which is also adopted in Tycoon decouples the lifetime and mobility of an entity from its type (orthogonal persistence [1], orthogonal mobility) and should be seen in contrast to systems and languages where the programmer has to tag persistent and mobile objects explicitly at creation-time. The latter systems provide a weaker notion of referential integrity since bindings from persistent (mobile) to volatile (immobile) objects are replaced at transaction-commit (migration-time) by bindings to a distinguished NIL entity or they may even become undefined.

We distinguish three categories of structured entities:

Data describes the persistent state of an information system by a collection of computational entities related through bindings. The structure (types) of the entities and their bindings are described by types. For example, the database schema of the *projectBudgetDB* at site *GroupManager* describes the signatures of the local persistent bindings between project and account entities stored at that site.

Code is a description of operation sequences that query and update volatile or persistent entities and bindings. Code can be expressed by means of imperative or declarative, high-level or low-level programs and scripts. For example, the impera-

tive script in Figure 3 is written at a rather high level of abstraction and involves static bindings to further code entities (*getExpenseDataFromUser*, *compileReport*) but also to mobile (*data*, *expenseReport*) and immobile (*projectBudgetDB*) data entities.

Threads are representations of code in the process of being executed. A thread is created by submitting a (non-parameterized) code fragment like the expense report script in Figure 3 and (persistent) data to an evaluator. Multiple threads executing the same code typically have different local bindings (*expenseReport*) but shared global bindings (*projectBudgetDB*).

The semantics of the thread evaluator can be defined inductively by rules that map thread states to thread states and that perform side-effects on data as described in more detail in [30]. A thread state subsumes bindings to the code fragments currently being executed and a dynamic environment that records the current bindings from names occurring in the code to local and global entities. In most imperative programming and query language implementations, thread states are represented as records that reference stacks of so-called "activation records", one for each function or query invocation (cf. Figure 1).

4. Thread Operations and Typing

Following Tycoon's add-on approach to data modeling [28], [29], the **migrate to** construct utilized in Figure 3 is not built into the core Tycoon system. Instead of this, migrating threads are provided by a hierarchy of library abstractions and Tycoon's extensible grammar [5] is exploited to provide the necessary syntactic layer to hide the underlying infrastructure from high-level workflow script programmers.

Tycoon's core thread functionality is provided by a library module that exports exactly the following exceptions, types and functions:

```
interface Thread export
(* — Types: *)
  T(R <:Ok) <:Ok
  State <:Ok
(* — Constants: *)
  error, abortion :Exception
  runningState, blockingState, suspendedState,
  terminatedState, exceptionState :State
(* — Queries: *)
  main() :T(Int)
  self() :T(Ok)
  running() :Array(T(Ok))
  state(R <:Ok  thread :T(R)) :State
(* — Thread management: *)
  new, fork, launch(R <:Ok  f(self :T(R)) :R) :Ok
  duplicate(R <:Ok  thread :T(R)) :T(R)
  suspend, run, abort, kill(R <:Ok  thread :T(R)) :Ok
```

```
  join(R <:Ok  thread :T(R)) :R
(* — Termination by exception: *)
  throw(R <:Ok  thread :T(R)  exc():Ok) :Ok
  catch(R<:Ok thread :T(R)  handler(exc():Ok):Ok):Ok
  joinCatch(R <:Ok  thread :T(R)) :R
(* — Synchronization primitives: *)
  atomic(R <:Ok  action() :R) :R
  sleep(timeout :Real) :Ok
end
```

In Tycoon, threads are *typed* first-class values that can be stored in variables, passed as parameters, embedded into (persistent) data structures and transmitted between network sites. For example, a variable of type *thread.T(Person)* can only hold a thread that evaluates code which returns a value of type *Person*. Tycoon's polymorphic typing makes it possible to define both, generic functions that work uniformly on threads with an arbitrary result type R (like *new, fork* or *suspend*), as well as user-defined functions that depend on a specific thread result type.

A Tycoon thread executes a function which is specified as an argument of the respective thread creation operation. For example, the following program fragment creates a new thread t which immediately start the execution of the function f that returns an integer value. It is required that a "thread script" function like f expects a thread parameter of the appropriate type, *T(Int)* in this case. At runtime, the *thread.fork* operation passes the newly created thread as an actual argument to the function to support type-safe reflexive thread operations.

```
let f(self :thread.T(Int)) :Int =   (* Thread-Skript *)
  begin ...
    thread.kill(self)   (* a reflexive operation on the current thread *)
    ...
    4711
  end
let t = thread.fork(f)   (* Creation of a new thread that executes f. *)
```

More precisely, the abstract data type *thread.T* is a type operator parametrized by the result type R of the thread function.

```
T(R <:Ok) <:Ok
```

By virtue of Tycoon's polymorphic type system, a thread result can be obtained in a type-safe way. For example, t is a value of type *thread.T(Int)* which implies that the result of the *thread.join* function applied to t returns an integer value.

```
let x :Int = thread.join(t)
```

At runtime, this statement blocks the execution of the current thread (*thread.self*) until the evaluation result of the function f is available. Thereupon execution of the waiting thread is resumed and the result of t is bound to the variable x. Mismatches

between the type (here *Int*) required by the given expression context and the thread result type are detected by the Tycoon compiler at compile-time

Without going into details, the functions *throw* and *catch* exploit the presence of higher-order functions in Tycoon to unify asynchronous signal handling with exception handling. For example, the function call *thread.abort(otherThread)* raises the exception *thread.abortion* in *otherThread*. In addition to critical sections (*thread.atomic*), several related modules export further synchronization primitives (semaphores, mutexes and condition variables) not shown here but discussed in detail in [37].

Building on these thread primitives, other Tycoon library modules add mechanisms to create a portable linear byte stream representation of a thread value, to establish a stream connection between network sites, to address network sites based on roles and logical identifiers, to atomically migrate running threads between network sites, etc. All these modules are loosely coupled and (by virtue of polymorphic typing and structural dynamic type checking across sites) strongly typed. This makes it possible to modify selected modules, for example, to experiment with different addressing and coordination protocols between migrating threads while ensuring a certain degree of overall system consistency.

5. Models for Distributed Activity-Oriented Systems

In this section we compare migrating threads with other programming models for distributed applications like remote procedure calls and (higher-order) distributed object management. We argue that in particular for activity-oriented tasks like workflow management it is desirable to add threads as mobile resources to these more traditional models. Our comparison is based on the uniform view of data, code and thread bindings introduced in Section 3.

In the following we restrict ourselves to distribution models that scale to heterogeneous, federated and widely distributed environments. Therefore, we do not discuss concurrent or database languages based on the notion of a single (persistent) address space, like distributed databases, distributed virtual memory or distributed persistent heaps, some of which already provide multi-threading capabilities. As argued in more detail in [26], we believe that the autonomy and heterogeneity requirements of larger-scale distributed systems are not addressed by these models.

5.1. Remote Procedure Call

Already in the introduction we quoted remote procedure calls as an example how to successfully generalize a well-established local programming concept to a distributed scenario. From a binding perspective, the core concept of RPCs are bindings from local to remote code entities which can be maintained for the duration of a session.

Typically, an activity is implemented by one thread of control at a single client side invoking subactivities at several server sites. A remote invocation either uses

a single thread at the server side which implicitly serializes requests from multiple clients or it creates a fresh volatile thread for each incoming client request (multi-threaded server).

The plain RPC model has the crucial disadvantage for activity-oriented programming, that the code describing a long-term activity has to be fragmented, distributed and installed statically at the relevant sites. Furthermore, with conventional RPC mechanisms [6], [36] it is very circumstantial to transmit recursive data structures.

5.2. Distributed Object Systems and Remote References

Distributed object management is viewed as a promising approach to build scalable distributed systems that are also capable of integrating legacy (database) systems by means of a unified object paradigm [24]. There are numerous proposals for specific object models like DSOM of IBM [14], DOM of GTE [24], Network Objects of Modula-3 [2] and future versions of Microsoft's OLE [31] and there are several related standardization efforts like CORBA of the OMG [12] and the OSF DCE/DME [36]. For a detailed feature analysis of these models see [35], [23].

From a binding perspective, distributed object models provide bindings to local and remote objects which aggregate data and code. Compared with the RPC model, the argument bindings for remote method invocation are generalized, since it is not only possible to transmit values but also references to other (local or remote) objects.

However, like in the RPC model there is no migration of behavior and object types have to be installed statically. As sketched in Figure 4 it is necessary to split activities into separate scripts stored as object methods at different sites. Moreover, long-term and recoverable activities are difficult to implement in today's distributed object models.

Figure 4 shows how the expense report activity described in Section 2 can be implemented by communicating distributed objects. Separate server objects describe the services available through the secretary, the project manager and the finance department. Activity control is centralized at the client (employee) site which is also responsible for exception handling at all stages of script execution.

The availability of remote references minimizes the data transfer during activity migration since each binding to a local entity at the originator site is replaced transparently on transmission by a remote reference. A remote reference identifies the originating address space and an entity local to it [18], [2], [4]. Subsequent read, write or execute access to the entity at the remote site transparently invokes a network communication with the originating site.

Conversion to remote references has clear, simple semantics since it preserves the original sharing. On the other hand it causes data fragmentation across address spaces, in particular it decreases the autonomy of migrating threads. Its usefulness depends on high site availability and low network latency.

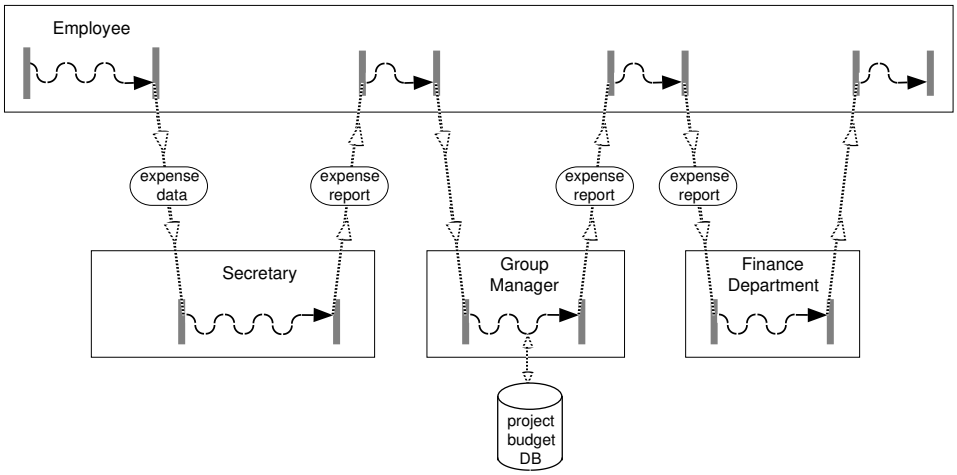

Figure 4. Implementing a distributed activity by communicating objects

5.3. Higher-Order Languages

Higher-order languages do not make a distinction between code and data bindings. In distributed higher-order languages like Obliq [4] it is therefore possible to transmit code and data uniformly between network sites. For example, by passing a function as an argument in an RPC, programmers can implement migration of behavior directly, as required by activity-oriented applications.

Moreover, since function abstraction dynamically aggregates state bindings, function migration already exhibits a limited mechanism to transmit partial evaluation states of activities across networks. For example, the Tycoon code in Figure 5 utilizes four functions (*step1* through *step4*) to encode the four evaluation states of the expense report script as valid at migration time. As expected, activity migrations are implemented by passing (the closures of) the above functions to higher-order RPCs [26] that simply execute their argument at the respective remote site. Such a *remote execution engine* can also establish bindings to resources available only at the receiver side by passing them as actual arguments to the function. For example, *groupManagerSite.execute* will execute *step3* with the *projectBudgetDB* as a local argument.

Note that in this approach iteration has to be expressed by recursion.

As illustrated by Figure 6, the resulting scenario is very close to our thread migration approach (see Figure 2). However, additional local thread functionality is required to implement concurrent activities and activities can only be transmitted

```
let rec step1 = fun() begin
  let data = getExpenseDataFromUser()
  let step2 = fun() begin
    let expenseReport=compileReport(data)
    if valid(expenseReport) then
      let step3=fun(projectBudgetDB :ProjectBudgetDB)
      begin
        update(projectBudgetDB expenseReport.total)
        let step4 = fun() begin
          transferMoney(expenseReport.total)
        end
        financeDepartmentSite.execute(step4)
      end
      groupManagerSite.execute(step3)
    else
      employeeSite.execute(step1)
    end
  end
  secretarySite.execute(step2)
end
employeeSite.execute(step1)
```

Figure 5. Encoding thread states with higher-order functions

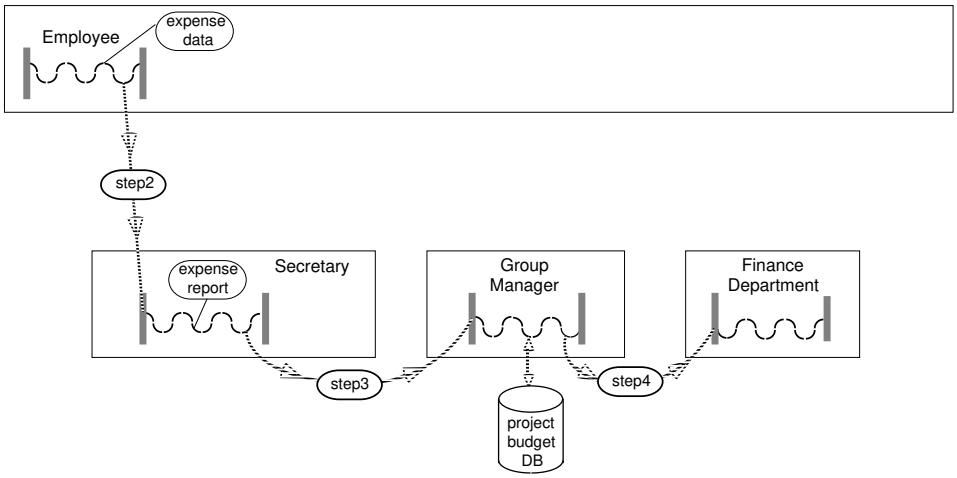

Figure 6. Emulating thread migration in a higher-order distributed language

in an execution state captured by an explicit function abstraction which destroys the block structure of imperative programs.

5.4. Migrating Objects

While entities in many distributed object models are immobile (they never leave their creation address space), some systems also support explicit object migration (Emerald [16], [18], SOS [40]). If an object leaves an address space, some local references are converted transparently into remote references. Conversely, if an object enters an address space, some remote references may be converted back into local references.

In addition to technical advantages (load balancing, reduced communication traffic, simplified system reconfiguration) object mobility provides migration of behavior and encapsulation of state as required for activity migration.

On the other hand, the state of a long-term activity represented by a migrating object has to be encoded explicitly as a (structured) object attribute. In particular, conditional, iterative and recursive state transitions lead to complicated state encodings. This explicit encoding should be seen in contrast to persistent threads where the runtime support implicitly and efficiently maintains the state of the long-term activity.

5.5. Migrating Threads and Network Agents

If one is interested in building systems where a large number of loosely coupled activities roams the network and makes heavy use of network resources, it is desirable to avoid a cumbersome encoding of activities and to have persistent and migrating first-class threads for straight-forward activity programming.

Such a complete unification of data, code and thread binding is achieved in Emerald [16], [17], [18] and in Tycoon. Emerald makes heavy use of reference semantics for these bindings. In particular, object state attributes and thread stack frames tend to be fragmented across multiple network nodes. Therefore, the main focus of Emerald is on applications in local area networks with high site availability and low network latency. As discussed in Section 6, Tycoon provides a spectrum of binding mechanisms to handle also other distribution scenarios, like autonomous network agents on global electronic marketplaces, a vision sketched in [45], [44].

Clearly, migrating threads are not a replacement for todays established distributed programming mechanisms, but they constitute a valuable programming abstraction that can be integrated smoothly into many distribution models.

6. Binding Techniques for Thread Resources

In this section we discuss the central issue in thread migration, namely how to transmit a thread state which represents a set of transitive bindings to data, code and other threads between address spaces. After a classification of thread resources we present tailored binding mechanisms provided in Tycoon for local and remote as well as for mobile and immobile resources.

6.1. Classification of Thread Resources

In order to determine appropriate binding mechanisms for resources that are involved in the migration of a thread, we first classify these resources based on their implementation technology:

Tycoon resources are all language entities (data, code, threads) which are defined entirely in Tycoon libraries or application code and which are therefore managed by the Tycoon runtime system.

External resources are conceptually pre-existing resources that are not under direct control of the Tycoon runtime system. These resources include files, communication channels, input or output devices, graphical elements (windows, buttons, menus) or programs (mail tool, word processor). External resources are created by external library code written in C or C++ which is available to Tycoon programmers via typed language gateways. In Section 6.6 we sketch a method that allows volatile external resources to "accompany" persistent and migrating threads through time and space.

In general, every atomic and structured Tycoon entity is persistent and **mobile**. If a component binding of a structured Tycoon entity (tuple, array, function, thread, ...) refers to an external resource, its persistence and mobility has to be ensured by explicit Tycoon code (see Section 6.6).

Some resources like platform-dependent software components or hardware devices are inherently **immobile**. In addition it may be necessary to regard some resources (for example, a huge database or a licensed piece of software) as immobile.

A further classification is based on the scope of resource definitions:

Local resources are defined in the inner lexical scope of a thread code definition and they are allocated dynamically for each thread (e.g., *data* and *expenseReport* in Figure 2).

Global resources are defined in the outer lexical scope of a thread code definition and they are potentially shared by multiple threads (e.g., the procedure definitions in Figure 2).

The following important cases of distributed resources require special binding methods:

Remote resources are available at a remote site only. For example, *projectBudgetDB* in Figure 2 is available exclusively at the group manager's site and requires a dynamic binding on thread migration from the secretary's site to the group manager's site.

Ubiquitous resources are available at all sites visited by a thread. Standard examples are the thread migration software itself (module *thread* in Section 4), operating system functions and program libraries that serve rather general purposes like GUI and network programming. Ubiquitous resources are usually stateless but an application-specific notion of resource equivalence frequently leads to a more general interpretation of ubiquity.

6.2. Shipping Mobile Resources

We explain the shipping of mobile resources using the following **migrate to** statement taken from Figure 3:

```
let data = ...
migrate to Secretary do
  let expenseReport = compileReport(data)
end
```

Assuming that the **migrate to** command is executed by a thread t running at site *Employee*, the migration of t is performed as follows. First, t is suspended at site *Employee*. Then, t is shipped to site *Secretary*. Finally, t is resumed at site *Secretary*. The notation **migrate to** ... **do** is realized by Tycoon's extensible syntax

[5] which translates directly into Tycoon library calls to perform these elementary thread manipulation and data communication steps.

As explained in Section 3, the thread t contains a continuation binding to the code still to be executed. This code fragment (***let*** *expenseReport = compileReport(data)*) in turn is represented as a function closure with bindings to all its free variables (*compileReport* and *data*). The free variables in the example above constitute a local thread resource (*data*, a data entity) and a global thread resource (*compileReport*, a code entity).

The basic semantics of thread shipping in Tycoon is a *deep copy* operation between address spaces. Therefore, the thread t at site *Secretary* contains direct local bindings to all mobile resources transitively reachable from t.

User-defined bulk data structures like lists, trees, database tables, etc., and also nested entities referring to code and threads are handled correctly by the Tycoon system. Programmers have to be aware that deep copying does not preserve sharing of mutable locations and may lead to high storage and communication costs.

In order to minimize the transitive referential closure of an entity, its representation must not include (indirect) references to entities which are irrelevant for its further use. A counterexample is a linked-list representation of function closures as found, for example, in Napier88 [33] where nested functions can unattendedly capture indirect references to bulk structures, as for instance the persistent root of the object store. This virtually prevents code and thread transmission and is also a severe obstacle to garbage collections. To avoid these problems, Tycoon uses flat function closures of minimal size determined by a static binding analysis during code generation.

6.3. Working with Immobile Resources

There are two alternatives to handle immobile resources like the database *projectBudgetDB* in Figure 3.

- Thread migration to the remote site (see Figure 2);
- Explicit communication with the remote site where the immobile resource is located.

In Tycoon the latter can be achieved by a remote procedure call (RPC) facility which is portable across different middleware architectures. As described in [26], there are two Tycoon RPC implementations based on ONC-RPC (also known as Sun-RPC) and BSD sockets respectively. A third implementation which utilizes DCE as its communication medium is in preparation.

6.4. Dynamic Binding to Remote Resources

In a thread script, only the static bindings to local resources can be checked at compile time against the local database schema. In a loosely coupled system where

remote sites can perform schema updates between thread script compilation time and thread migration time, it is necessary to perform *dynamic* bindings to remote resources (e.g. *projectBudgetDB* in Figure 3). In order to locate a remote resource within the destination address space, an ubiquitous *name service* like the network object import mechanism in Modula-3 [2] or Obliq [4] can be used to identify resources based on a string value. This requires a dynamic type check between the actual resource type and the types used in script programming and bears the risk that a type error is detected after migration only.

As an alternative, the types of the remote resources can be attached already to the type specification of a remote *migration engine*. A migration engine is an RPC server that accepts a thread, binds it dynamically to local site resources and then resumes the thread. In Tycoon, a remote migration engine is typed based on the signatures of its resources. Therefore, mismatches between remote types and the types used in thread scripts are detected already at server binding time. Thus type errors are limited to the departure site and can only happen once per connection.

Syntactically, the name and the type of each remote resource is listed in the **with remote** clause of a migrate statement. The scope of these identifiers is restricted to the block enclosed by the keywords **do** and **end** (see also Figure 3):

```
migrate to GroupManager with remote
  projectBudgetDB :ProjectBudgetDB
do
   update(projectBudgetDB, expenseReport.total)
end
```

To summarize, Tycoon thread scripts are fully statically type checked at their originating site. Whenever two sites establish a network connection (which may lead to a large number of thread migrations or RPC calls), the resource type definitions of the two sites are verified to be (structurally) compatible.

6.5. Access to Ubiquitous Resources

Experience with non-trivial activity-oriented applications shows that in order to avoid excessive communication or storage costs on thread migration, system support is needed to handle bindings to ubiquitous resources present at each node visited by a migrating thread on the network.

From the programmer's viewpoint, it suffices to mark ubiquitous resources with a call to the function *register* of the module *ubiquitous*.

```
ubiquitous.register(windowManager)
ubiquitous.register(list)
ubiquitous.register(database)
```

In the example above, the Tycoon modules *windowManager*, *list* (polymorphic lists) and an application-specific module *database* are tagged as ubiquitous resources. The argument of the *register* function can be an arbitrary Tycoon entity;

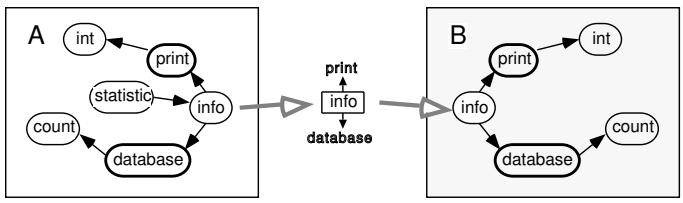

Figure 7. Rebinding ubiquitous resources by means of dynamic linking

it is not constrained to coarse-grained module values. This is a generalization of the limited dynamic linking capabilities of other distributed programming systems (see e.g., Emerald [16], SOS [40], [39], FACILE [20] and Java [9]).

On transmission between address spaces, bindings to entities explicitly marked as ubiquitous are replaced by symbolic references. At the receiver site, these symbolic references are replaced by corresponding local bindings. In principle, a dynamic type check has to be performed during this *dynamic linking* to ensure that the type of the remote resource matches the type of the resource at the originator site. In the Tycoon implementation, this expensive repeated operation can be avoided by global identification mechanisms (for details, see [26]).

The dynamic linking of ubiquitous resources is most useful for immutable values (code, literals, external bindings) replicated over the network. However, there are also some situations where a migrating thread has to bind to mutable repositories at multiple network sites. For example, a migrating thread might operate on multiple databases, file systems or ftp servers or produce side-effects on multiple screens, fax machines, etc. while roaming the network.

The following Tycoon code shows the definition of the function *info* in the module *statistics* which is bound to the ubiquitous modules *database* and *print*.

```
module statistic import database print
export let info() = print.int(database.count())
end;
```

The left hand side of Figure 7 illustrates the bindings within the address space A that contains the module *statistics*. If a thread that uses solely the function *info* is transferred to another address space B, the deep copy operation stops at the ubiquitous resources *database* and *print*.

6.6. Recreation of Volatile Resources

Volatile resources like window handles and bindings to C data structures outside of the Tycoon persistent store can be registered with the module *volatile* in order to

make them pseudo-persistent. Such volatiles are recreated automatically following a system restart or a migration operation and they are destructed automatically preceding shutdown, rollback or migration operations.

The order in which volatiles are registered is significant, because typically there exist bindings from "younger" to "older" volatiles. If programmers register volatiles in creation order, the automatic recreation operations follow the original creation sequence and destruction happens in reverse order. A similar mechanism is embodied in the SOS system, where resource dependencies are defined by naming *pre-requisite objects* [40] which have to be recreated before their dependent objects.

Each Tycoon address space contains a global data structure referencing all volatiles under the control of the module *volatile*. The elements of this data structure never migrate; instead they are recreated automatically at the receiver site. Volatiles interact with the Tycoon garbage collector through "weak references" [13].

7. Implementing Migrating Persistent Threads

The previous sections described migrating threads from an application programmer's point of view. We now turn to implementation aspects of thread migration, namely storage and data representation (Section 7.1), communication protocols (Section 7.2) and transaction management (Section 7.3).

7.1. A Mobile and Persistent Thread State Representation

Figure 1 shows the most important constituents of the execution (evaluation) state of a local, volatile thread. Such a thread is represented by a data structure that captures the states of all relevant registers of the CPU. These registers, in particular the stack pointer and the program counter, constitute transitive references to all further components of the evaluation state. Occasionally, the transitive referential closures of several threads overlap. In this simple way, local volatile threads can share program resources of all kinds like code, variables, files and communication channels.

From the programmer's point of view, Tycoon threads behave just like normal volatile threads in other programming languages. However, Tycoon thread representations are integrated with the global transaction and persistence mechanisms of Tycoon. Figure 8 illustrates the concrete representation of a Tycoon thread by persistent store objects. The object in the upper left corner constitutes the actual thread value (a handle). It contains the thread state, a reference to the thread stack and the result. Whereas the former is a persistent entry, the latter are valid only in certain thread states. In the following, the set of possible thread states *state* and the respective transitions are explained (see Figure 9.

running: Expression evaluation takes place in this state. Each executing thread has its own persistent stack object referenced by the field *stack*. The first slots

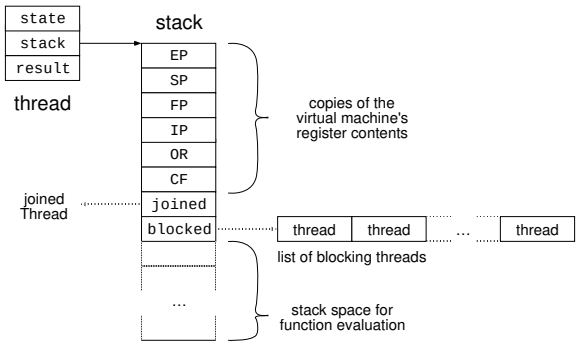

Figure 8. Object store representation of a persistent and mobile Tycoon thread

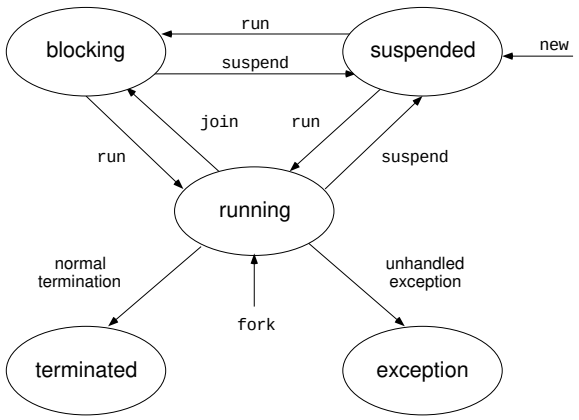

Figure 9. Persistent thread states

of the stack object contain the abstract machine's register values (*IP* ... *CF*) and some extra administrative data. By storing these and the stack contents in a single object, allocation and release of registers and stacks is simplified.

suspended: In this state a thread has been stopped explicitly by a call to *thread.suspend*. This does not change its storage structure which also applies to the state *blocked*.

blocking: A thread that waits for the result of another thread (due to a call to *thread.join*) is blocked until the required value becomes available. A blocking thread can also be suspended.

terminated: When a thread has completed the evaluation of its function, the function result is stored in the field *result* and propagated to all threads that have a pending *thread.join* on this thread. Then the reference to the *stack* is removed. Thus, a terminated thread occupies only very little storage space, but still provides its result to other threads that still have a reference to this thread and may execute *thread.join* at a later point in time.

exception: This state is reached when a dynamic exception is raised during evaluation that propagates through all active function frames of a thread and that is not handled by a matching exception handler installed by a **try end** block. The exception package is stored in the field *result* and is propagated to all threads that execute *thread.join* on this thread. The field *stack* is treated as described for the state *terminated*.

The evaluation of a Tycoon thread does not use the persistent representation of its registers directly to avoid expensive persistent store accesses in the virtual machine. Instead of this, there is also a volatile representation for all frequently used thread components such as registers and thread stacks which are mapped to main memory by a special object store operation *tsp_openReadWriteLock* [27]. Volatile and persistent thread representations are synchronized prior to the following events: thread state transition, global garbage collection, transaction commit and transaction restart.

In case of a global transaction commit, all threads in a *running* state are captured in a global list which constitutes the root of the persistent object store. This design decision leads to a simple and clean storage management strategy:

- All objects (in particular passive threads) that are transitively reachable by an active thread are recognized by the garbage collection as being "alive" and are stored persistently.
- All other objects are garbage collected since there is no further activity that can access them directly or indirectly.

We have designed the thread representation carefully to reduce references between threads to a minimum. This is not only necessary to avoid persistent storage leaks, but also to preserve the autonomy of threads which is important for their mobility.

7.2. Typed Thread Migration Gates

In this section we describe how we implement migration of threads using local thread manipulation operations and on top of a standard RPC communication layer.

A simple "passive" thread migration takes place when a thread is passed as a direct or indirect RPC-argument by another active thread. Workflows and agents typically require an "active" thread migration since they initiate their migration themselves. This requires some non-trivial implementation steps, because the transmission of a migrating thread cannot be controlled completely by itself since it is necessary to distinguish two evaluation states: the state of the thread to be shipped and resumed at the receiver site as well as the state of the thread that controls the communication steps. Moreover type-safe bindings to local resources at the receiver site require special arrangements.

The module *gate* implements active thread migration. Since this module is implemented completely in Tycoon, is is possible to describe active thread migration in the remainder of this subsection at a high level of abstraction by explaining the Tycoon source code, line by line. The module *gate* is based on Tycoon's type-safe and type-complete RPC-mechanism which is described in [26] and [25].

First of all, the module *gate* defines a type operator *Parameter(Data)* which is parametrized by the type *Data* of resources which will be bound dynamically at the receiver site.

```
Let Parameter(Data <:Ok) = Tuple
  var data :optional.T(Data)
  agent :thread.T(Ok)
end
```

The implementation of this type operator is a tuple that aggregates the migrating thread with a placeholder for the resources that will be supplied in the target store.

The main type of the module (called *gate.T*) are remote *gates* to which threads can migrate in order to enter remote object stores:

```
Let T(Data <:Ok) <:Service = Tuple
  transfer(parameter :Parameter(Data)) :Ok
end
```

A value of type *gate.T(Data)* is an RPC service which provides a single function. This function (*transfer*) is invoked by *gate.migrateTo* to transmit packages of the matching type *Parameter(Data)*.

A typical application of *gate.migrateTo* looks as follows:

```
let groupManagerSite = client.bind(...  :T(ProjectBudgetDB))
...
let projectBudgetDB = gate.migrateTo(groupManagerSite)
...
```

First, the remote gate *groupManagerSite* is made accessible to the current thread by means of an RPC client binding. When the thread "passes" this gate with the *migrateTo* operation, a value that is determined by the respective RPC service is returned and enters the scope of the thread function which is now being executed at the remote site.

The type-safe implementation of *gate.migrateTo* looks as follows:

```
let migrateTo(Data <:Ok  gate :T(Data)) :Data =
  begin
    let parameter = tuple
      let var data = optional.nil(:Data)
      let agent = thread.self()
    end
    thread.launch(fun(self :thread.T(Ok))
      begin
        gate.transfer(parameter)
        thread.kill(parameter.agent)
      end)
    optional.value(parameter.data)
  end
```

First, a package is created which aggregates the current thread and a null value (placeholder) of type *Data* which will be updated by another function at the receiver site. Second, a fresh thread is launched that carries out the migration operations on the now "passive" thread *parameter.agent*. Finally, the value stored in the field *parameter.date* is returned to the caller (which is now being executed at the remote site).

The transmission of the package *parameter* is executed by a short-lived local helper thread which is created (and started) by the thread primitive *thread.launch*. The latter acts like *thread.fork*, but atomically also suspends the calling thread. The local helper thread executes the anonymous function that begins with the keyword **fun**. First, it calls the RPC *gate.transfer* which transmits the package. Then it kills the local representation of the already migrated thread. After this, the helper thread terminates and as there are no references to it, the garbage collection will release its resources.

The statement *optional.value(parameter.data)* which yields the function result of *migrateTo* is executed by a copy of the migrated thread on the recipient site. This results from the implementation of the RPC service which constitutes the migration gate.

At the server site, the function *gate.new* creates a migration gate and registers it dynamically as an RPC service.

```
let new(dispatcher :server.T  ...  data() :Data) :sid.T(T(Data)) =
  begin
    let gate = tuple
      let transfer(parameter :Parameter(Data)) =
```

```
      begin
        parameter.data := optional.new(data())
        thread.run(parameter.agent)
      end
    end
    server.register(dispatcher  ...  gate)
  end
```

The parameter *dispatcher* and further parameters not shown here are simply passed on to the function *server.register* which registers a service *gate* at the *dispatcher*. The type *Data* indirectly determines the type (*T(Data)*) of the RPC service *gate* which is constructed dynamically by the first binding in the function. The function *transfer* offered by the RPC service *gate* takes a parameter of type *Parameter(Data)* that matches the packets generated by the *migrateTo* function described above.

Note that access to the function *data* stems from a dynamic binding, i.e. parameter passing. This dynamic binding and the assignment to *parameter.data* implement the transfer of a local value into the scope of a migrating thread in a type-safe way. Moreover, it should be noted that Tycoon provides higher-order functions such that the parameterless function *data* can provide arbitrary objects from within the whole object store.

The examples in this subsection also demonstrate that Tycoon's polymorphic type system is sufficiently expressive to describe the migration of threads with strongly-typed bindings to remote resources without resorting to low-level, type-unsafe language primitives.

7.3. Transactional Thread Migrations

The preferred application domain for migrating threads is a loosely coupled environment with highly autonomous subsystems. In such an environment, consistency maintenance is limited to local data stores only. If integrity constraints do not span multiple systems, there is no need for large spheres of control [11] which have to protect concurrent distributed activities against unwanted interference.

Migrating activities that are implemented by threads should refrain, as far as possible, from complicating these simple conditions. Therefore we advocate binding techniques that support the autonomy of migrating threads and in consequence the autonomy of visited sites. The problematic nature of distributed transactions is avoided wherever possible.

On the other hand, a thread migration itself may lead to concurrency-control and recovery problems. In particular, the departure and the destination site should establish a consistent view of the outcome of each migration step. Otherwise, a thread instance could be duplicated or be lost. As a consequence, a migration should be performed as a distributed transaction. For this purpose, the well-known two-phase commit technique [10], [22], [21], [15] is sufficient.

During migration, the client of a gate takes the role of a transaction coordinator and of a transaction participant. The gate server acts as a participant only.

It is important to suspend the reactivation of the migrating thread until the transaction is completed. Otherwise, it could happen that some actions must be undone in case of a transaction failure.

Besides the departure and the destination site of a migrating thread there can be more sites that are indirectly affected by the transaction, because the migration itself might be an application-relevant information. For example, it might be interesting for a manager to know at which site a certain workflow is situated or a user might want to know how far in physical terms an agent has already traveled. As Tycoon does neither restrict distributed programming to agent style (like Telescript [8]) nor to a pure client/server style, there may also be "tethered agents" that are more or less under remote control by a third-party site while they are roaming the network. It is straight-forward to extend the simple migration transaction scheme by a set of further 2PC participants to suit such scenarios.

8. Conclusion

We propose to enrich today's distributed programming models with higher-order concepts like first-class persistent and migrating threads and functions to provide better support for long-lived distributed activities like process modeling [7], workflow management and network agent programming [45], [44], [43].

At the system level, thread migration can be supported uniformly by techniques that are well-established in the persistent programming language world: stream representations of tagged object-graphs, portable code formats, function closures and polymorphic structural type checking. However, our experience with non-trivial activity-oriented applications shows that in order to avoid excessive communication or storage costs, additional system and language support has to be provided to support bindings to remote and ubiquitous resources.

More work is required to scale thread synchronization and coordination to a distributed, persistent scenario and to investigate the relationship between threads and transactions [3].

References

1. M.P. Atkinson and P. Bunemann. Types and persistence in database programming languages. *ACM Computing Surveys*, 19(2), June 1987.
2. A. Birell, G. Nelson, S. Owicki, and E. Wobber. Network objects. In *14th ACM Symposium on Operating System Principles*, pages 217–230, June 1993.
3. Y. Breibart, A. Deacon, H.-J. Schek, A. Sheth, and G. Weikum. Merging application-centric and data-centric approaches to support transaction-oriented multi-system workflows. *ACM SIGMOD Record*, 12(3):23–30, September 1993.
4. L. Cardelli. Obliq: A language with distributed scope. Technical report, Digital Equipment Corporation, Systems Research Center, Palo Alto, California, June 1994.

5. L. Cardelli, F. Matthes, and M. Abadi. Extensible grammars for language specialization. In C. Beeri, A. Ohori, and D.E. Shasha, editors, *Proceedings of the Fourth International Workshop on Database Programming Languages, Manhatten, New York*, Workshops in Computing, pages 11–31. Springer-Verlag, February 1994.
6. J.R. Corbin. *The Art of Distributed Applications.* Sun Technical Reference Library. Springer-Verlag, 1991.
7. B. Curtis, M.I. Kellner, and J. Over. Process modelling. *Communications of the ACM*, 35(9), September 1992.
8. Telescript programming guide, version 0.5 (alpha). http://www.genmagic.com/Telescript/TDE/TDEDOCS_HTML/developer.html, October 1995.
9. J. Gosling and H. McGilton. The java language environment – A whitepaper. Technical report, Sun Microsystems, October 1995.
10. J. Gray. Notes on database operating systems. In *Operating Systems – An Advanced Course*, volume 60 of *Lecture Notes in Computer Science*. Springer-Verlag, 1978.
11. J. Gray and A. Reuter. *Transaction Processing – Concepts and Techniques.* The Morgan Kaufmann Series in Data Management Systems. Morgan Kaufmann Publishers, 1993.
12. Object Management Group. The common object request broker: Architecture and specification. Document 91.12.1, Rev. 1.1, OMG, December 1991.
13. J. Horning, P. Kalsow, J. McJones, and G. Nelson. Some useful Modula-3 interfaces. Technical report, Digital Equipment Corporation, Systems Research Center, Palo Alto, California, December 1993.
14. IBM Corporation, Publication No. SR28-5570. *Object-Oriented Programming using SOM and DSOM*, August 1994.
15. ISO / IEC JTC1 / SC21 / IS 9805. *Information Technology - Open Systems Interconnection - Protocol Specification for the Commitment, Concurrency and Recovery Service Element*, 1990.
16. E. Jul. *Object Mobility in a Distributed Object-Oriented System.* PhD thesis, Department of Computer Science, University of Washington, Seattle, Washington, 1988.
17. E. Jul. Migration of light-weight processes in Emerald. *Operation Systems Technical Comittee Newsletter*, 3(1):25–30, 1989.
18. E. Jul, H. Levy, N. Hutchinson, and A. Black. Fine-grained mobility in the Emerald system. *ACM Transactions of Computer Systems*, 6(1):109–133, February 1988.
19. D. Keppel. Tools and techniques for building fast portable threads packages. Technical Report 93-05-06, University of Washington, Department of Computer Science and Engineering, University of Washington, Seattle, May 1993.
20. F. Knabe. *Language Support for Mobile Agents.* PhD thesis, School of Computer Science, Carnegie-Mellon University, December 1995.
21. B. W. Lampson. Atomic transactions. In *Distributed Systems–Architecture and Implementation*, volume 105 of *Lecture Notes in Computer Science*, pages 246–265. Springer-Verlag, 1981.
22. B. Lindsay, P. Selinger, C. Galtieri, J. Gray, R. Lorie, T. Price, F. Putzolo, I. Traiger, and B. Wade. Notes on distributed databases. Technical Report TR RJ2571, IBM Almaden Research Laboratories, San Jose, CA, 1979.
23. F. Manola and S. Heiler. A "RISC" object model for object system interoperation: Concepts and applications. Technical Report TR-0231-08-93-165, GTE laboratories Inc., Waltham, MA (USA), August 1993.
24. F. Manola, S. Heiler, D. Georgakopoulos, M. Hornick, and M. Brodie. Distributed object management. *International Journal of Intelligent and Cooperative Information Systems*, 1(1), March 1992.
25. B. Mathiske. *Mobility in Persistent Object Systems.* PhD thesis, Fachbereich Informatik, Universität Hamburg, Germany, May 1996. (in German).
26. B. Mathiske, F. Matthes, and J.W. Schmidt. Scaling database languages to higher-order distributed programming. In *Proceedings of the Fifth International Workshop on Database Programming Languages, Gubbio, Italy.* Springer-Verlag, September 1995. (Also appeared as TR FIDE/95/137).

27. F. Matthes, R. Müller, and J.W. Schmidt. Towards a unified model of untyped object stores: Experience with the Tycoon store protocol. In M.P. Atkinson, editor, *Fully Integrated Data Environments*. Springer-Verlag (to appear), 1995.
28. F. Matthes and J.W. Schmidt. Bulk types: Built-in or add-on? In *Database Programming Languages: Bulk Types and Persistent Data*. Morgan Kaufmann Publishers, September 1991.
29. F. Matthes and J.W. Schmidt. Definition of the Tycoon language TL – a preliminary report. Informatik Fachbericht FBI-HH-B-160/92, Fachbereich Informatik, Universität Hamburg, Germany, November 1992.
30. F. Matthes and J.W. Schmidt. Persistent threads. In *Proceedings of the Twentieth International Conference on Very Large Data Bases, VLDB*, pages 403–414, Santiago, Chile, September 1994.
31. Microsoft Corporation. *Microsoft Office Developer's Kit*, 1994.
32. M. Morrison, P.T. Atkinson, A.L. Brown, and A. Dearle. Bindings in persistent programming languages. In *SIGPLAN Notices*, volume 23, pages 27–34, April 1988.
33. R. Morrison, A.L. Brown, R. Connor, and A. Dearle. The Napier88 reference manual. PPRR 77-89, Universities of Glasgow and St Andrews, 1989.
34. G. Nelson, editor. *Systems programming with Modula-3*. Series in innovative technology. Prentice Hall, Englewood Cliffs, New Jersey, 1991.
35. J. Nicol, T. Wilkes, and F. Manola. Object orientation in heterogeneous distributed computing systems. *Special Issue on Heterogeneous Processing*, June 1993.
36. OSF. *OSF DCE Administration Guide – Core Components*. Prentice Hall, Englewood Cliffs, New Jersey, 1993.
37. A. Piellusch. Synchronization of long-lived activities. Master's thesis, Fachbereich Informatik, Universität Hamburg, Germany, June 1996. (in German).
38. Portable operating system interface for computer environments (POSIX). Federal information processing standards publication NBS-FIPS-PUB-151-1, National Bureau of Standards, 1990.
39. M. Shapiro. Flexible bindings for fine-grain and fragmented objects in distributed systems. Rapport de Recherche 2007, INRIA, Domaine de Voluceau, Rocquencourt 78153 Le Chesnay Cedex, France, August 1993.
40. M. Shapiro, P. Gautron, and L. Mosseri. Persistence and migration for C++ objects. In *In Proceedings of the European Conference on Object Oriented Programming, Nottingham, GB*, July 1989.
41. C. Strachey, editor. *Fundamental concepts in programming languages*. Oxford University Press, Oxford, 1967.
42. Sun Microsystems. *Sun OS Reference Manual*, May 1988.
43. B. Thomsen, L. Leth, F. Knabe, and P.Y. Chevalier. Mobile agents. Technical Report ECRC-95-21, European Computer-Industry Research Centre, Munich, Germany, June 1995.
44. P. Wayner. Agents away. *BYTE*, pages 113–118, May 1994.
45. J.E White. Telescript technology: The foundation for the electronic marketplace. White paper, General Magic Inc., Mountain View, California, USA, 1994.