



Federated Method Invocation with Exertions

Michael Sobolewski

Computer Science, Texas Tech University
SORCER Research Group, <http://sorcer.cs.ttu.edu>
sobol@cs.ttu.edu

Abstract. Six generations of RPC systems can be distinguished including Federated Method Invocation (FMI) presented in this paper. Some of them—CORBA, Java RMI, and Web/OGSA services—support distributed objects. However, creating object wrappers implementing remote interfaces doesn't have a great deal to do with object-oriented distributed programming. Distributed objects developed that way are usually ill-structured with missing core object-oriented traits: encapsulation, instantiation, inheritance, and network-centric messaging by ignoring the real nature of networking. A distributed system is not just a collection of distributed objects—it's the network of dynamic objects. In particular, the object wrapping approach does not help to cope with network-centric messaging, invocation latency, object discovery, dynamic object federations, fault detection, recovery, partial failure, etc. The Jini™ architecture does not hide the network; it allows the programmer to deal with the network reality: leases for network resources, distributed events, transactions, and discovery/join protocols to form service federations. An exertion-based architecture presented in this paper implements FMI to support service-oriented metaprogramming. The new triple Command pattern architecture presented in this paper uses Jini service management for managing the network of FMI objects.

1 Introduction

Socket-based communication forces us to design distributed applications using a read/write (input/output) interface, which is not how we generally design non-distributed applications based on procedure call (request/response) communication. In 1983, Birrell and Nelson devised remote procedure call (RPC) [2], a mechanism to allow programs to call procedures on other hosts. So far, six RPC generations can be distinguished:

1. First generation RPCs [2]—Sun RPC (ONC RPC) and DCE RPC, which are language, architecture, and OS independent;
2. Second generation RPCs—CORBA [20] and Microsoft DCOM-ORPC, which add distributed object support;
3. Third generation RPC—Java RMI [17] is conceptually similar to the second generation but supports the semantics of object invocation in different address spaces that are built for Java only. RMI fits cleanly into the language with no need for standardized data representation, external interface definition language, and with behavioral transfer that allows remote objects to perform operations that are determined at runtime;

4. Fourth generation RPCs—next generation of Java RMI, Jini Extensible Remote Invocation (Jini ERI) with dynamic proxies, smart proxies, network security, and with dependency injection defining exporters, end points, and security;
5. Fifth generation RPCs—Web/OGSA Services RPC [16, 27] and the XML movement including Microsoft WCF/.NET;
6. Sixth generation RPC—Federated Method Invocation (FMI), presented in this paper, allows for concurrent invocations on multiple federating hosts (virtual metacomputer) in the SORCER environment [25].

All the RPC generations are based on a form of service-oriented architecture (SOA) discussed in Section 2. However, CORBA, RMI, and Web/OGSA services are in fact object-oriented wrappers of network interfaces that hide distribution and ignore the real nature of network through classical abstractions of object-oriented programming using existing network technologies. The fact that object-oriented languages are used to create these object wrappers doesn't mean that developed distributed objects have a great deal to do with object-oriented distributed programming. For example, CORBA defines many services, and implementing them using distributed objects does not make them well structured with core object-oriented traits: encapsulation, instantiation, inheritance, and network-centric messaging. Similarly in RMI, marking objects with the `Remote` interface does not help to cope with network-centric messaging, object discovery, dynamic federation, fault detection, recovery, partial failure, etc.

Programmers use abstractions all the time. The source code written in programming language is an abstraction of the machine language. From machine language to object-oriented programming, layers of abstractions have accumulated like geological strata. Every generation of SW architects and programmers uses it's era's programming languages and tools to build programs of the next generation. Each architecture and programming language used reflects a relevant abstraction, and usually the type and quality of the abstraction implies the complexity of problems we are able to solve. For example, a procedural language provides an abstraction of an underlying machine language. Building on the object-oriented distributed paradigm is the Federated Service Object-Oriented (FSOO) paradigm exemplified by the Jini architecture [12] in which the network objects come together on the fly to play their predefined roles. In the Service-ORiented Computing EnviRonmet (SORCER) developed at Texas Tech University [25], a service provider is a remote object that accepts network requests — called *exertions*—from service requestors to execute an item of work. While conventional objects encapsulate *data* and *operations*, exertions encapsulate *data*, *operations*, and *control strategy*. An exertion can federate on multiple hosts according to its *control strategy*.

The SORCER metacomputing environment adds an entirely new layer of abstraction to the practice of grid computing—*exertion-oriented programming*. The exertion-oriented programming makes a positive difference in service-oriented programming primarily through a new metaprogramming abstraction as experienced in many grid-computing projects including projects deployed at GE Global Research Center, GE Aviation, Air Force Research Lab, and SORCER Lab [19, 8, 9, 14, 22].

An exertion submitted to any provider in SORCER dynamically bounds to all relevant and currently available service providers on the network. The providers that

dynamically participate in this invocation are collectively called an *exertion federation*. This federation is also called a *virtual metacomputer* since federating services are located on multiple physical compute nodes held together by the FSOO infrastructure so that, to the individual exertion requestor, it looks and acts like a single computer.

The SORCER environment provides the means to create interactive FSOO programs [22] and execute them using the SORCER runtime infrastructure presented in Section 3. Exertions can be created using interactive user interfaces downloaded on the fly from service providers. Using these interfaces, the user can execute and monitor the execution of exertions within the FSOO metacomputer. The exertions can be persisted for later reuse, allowing the user to quickly create new applications or programs on the fly in terms of existing, usually persisted, exertions.

SORCER is based on the evolution of concepts and lessons learned in the FIPER project [5,21], a \$21.5 million program founded by NIST (National Institute of Standards and Technology). Academic research on exertion-oriented programming has been established at the SORCER Laboratory, TTU, [25] where 23 SORCER related research studies have been investigated so far [26]. The current version of FMI used in SORCER is described in this paper.

The paper is organized as follows. Section 2 provides a brief description of a service oriented architecture with a related discussion of distribution transparency; Section 3 describes the SORCER methodology; Section 4 presents federated method invocation; Section 5 provides concluding remarks.

2 SOA and Distribution Transparency

Various definitions of a Service-Oriented Architecture (SOA) leave a lot of room for interpretation. In general terms, SOA is a software architecture using loosely coupled software services that integrates them into a distributed computing system by means of service-oriented programming. Service providers in the SOA environment are made available as independent service components that can be accessed without a priori knowledge of their underlying platform or implementation. While the client-server architecture separates a client from a server, SOA introduces a third component, a service registry. In SOA, the client is referred to as a service requestor and the server as a service provider. The provider is responsible for deploying a service on the network, publishing its service to one or more registries, and allowing requestors to bind and execute the service. Providers advertise their availability on the network; registries intercept these announcements and add published services. The requestor looks up a service by sending queries to registries and making selections from the available services. Requestors and providers can use discovery and join protocols to locate registries and then publish or acquire services on the network. We can distinguish the *service object-oriented architectures* (SOOA), where providers are network objects accepting remote invocations, from *service protocol oriented architectures* (SPOA), where a communication protocol is fixed and known beforehand by the provider and requestor. Based on that protocol and a service description obtained from the service registry, the requestor can bind to the service provider by creating a proxy used for remote communication over the fixed protocol. In SPOA a service is usually identified by a name. If a service provider registers its

service description by name, the requestors have to know the name of the service beforehand.

In SOOA, a proxy—an object implementing the same service interfaces as its service provider—is registered with the registries and it is always ready for use by requestors. Thus, in SOOA, the service provider publishes the proxy as the active surrogate object with a codebase annotation, e.g., URLs to the code defining proxy behavior (RMI and Jini ERI). In SPOA, by contrast, a passive service description is registered (e.g., an XML document in WSDL for Web/OGSA services, or an interface description in IDL for CORBA); the requestor then has to generate the proxy (a stub forwarding calls to a provider) based on a service description and the fixed communication protocol (e.g., SOAP in Web/OGSA services, IIOP in CORBA). This is referred to as a bind operation. The binding operation is not required in SOOA since the requestor holds the active surrogate object obtained from the registry.

Web services and OGSA services cannot change the communication protocol between requestors and providers while the SOOA approach is protocol neutral [30]. In SOOA, how an object proxy communicates with a provider is established by the contract between the provider and its published proxy and defined by the provider implementation. The proxy's requestor does not need to know who implements the interface or how it is implemented. So-called smart proxies (Jini ERI) grant access to local and remote resources; they can also communicate with multiple providers on the network regardless of who originally registered the proxy. Thus, separate providers on the network can implement different parts of the smart proxy interface. Communication protocols may also vary, and a single smart proxy can also talk over multiple protocols including application specific protocols.

SPOA and SOOA differ in their method of discovering the service registry. SORCER uses dynamic discovery protocols to locate available registries (lookup services) as defined in the Jini architecture [12]. Neither the requestor who is looking up a proxy by its interfaces nor the provider registering a proxy needs to know specific locations. In SPOA, however, the requestor and provider usually do need to know the explicit location of the service registry—e.g., the IP address of an ONC/RPC portmapper, a URL for RMI registry, a URL for UDDI registry, an IP address of a COS Name Server—to open a static connection and find or register a service. In deployment of Web and OGSA services, a UDDI registry is sometimes even omitted (WSDL descriptions are shared via files outside of the system); in SOOA, lookup services are mandatory due to the dynamic nature of objects identified by service types. Interactions in SPOA are more like client-server connections (e.g., HTTP, SOAP, IIOP), in many cases with no need to use service registries at all.

Let's emphasize the major distinction between SOOA and SPOA: in SOOA, a proxy is created and always owned by the service provider, but in SPOA, the requestor creates and owns a proxy which has to meet the requirements of the protocol that the provider and requestor agreed upon a priori. Thus, in SPOA the protocol is always a generic one, reduced to a common denominator—one size fits all—that leads to inefficient network communication in some cases. In SOOA, each provider can decide on the most efficient protocol(s) needed for a particular distributed application.

Service providers in SOOA can be considered as independent network objects finding each other via service registries and communicating through message passing. A collection of these objects sending and receiving messages—the only way these

objects communicate with one another—looks very much like a service object-oriented distributed system.

Do you remember the eight fallacies of network computing? [4] We cannot just take an object-oriented program developed without distribution in mind and make it a distributed system, ignoring the unpredictable network behavior. Most RPC systems, except Jini, hide the network behavior and try to transform local communication into remote communication by creating distribution transparency based on a local assumption of what the network might be. However every single distributed object cannot do that in a uniform way as the network is a heterogeneous distributed system and cannot be represented completely within a single entity.

The network is dynamic, can't be constant, and introduces latency for remote invocations. Network latency also depends on potential failure handling and recovery mechanisms so we cannot assume that a local invocation is similar to remote invocation. Thus complete transparency distribution—by making calls on distributed objects as though they were local—is impossible to achieve in practice. The distribution is not just an object-oriented implementation of a single type of distributed object; it's a metasystemic issue in object-oriented distributed programming.

Exertion-based programming is introduced to handle the metasystemic distribution in SORCER by using indirect remote method invocation with no service provider explicitly specified in the network request (exertion). Specific infrastructure objects support exertion-oriented programming combined with FMI. That infrastructure defines SORCER's distributed object modularity, extensibility, and reuse of service-oriented components consistent with the relevant metacomputing granularity and dependency injection—key features of object-oriented distributed programming that are usually missing in SPOA programming environments.

3 Federated Service Object-oriented Computing Environment: SORCER

SORCER is a federated service-to-service (S2S) metacomputing environment that treats service providers as network objects with well-defined semantics of a federated service object-oriented architecture (FSOOA). It is based on Jini semantics of services [12] in the network and Jini programming model with explicit leases, distributed events, transactions, and discovery/join protocols. While Jini focuses on service management in a networked environment, SORCER is focused on exertion-oriented programming and the execution environment for exertions. SORCER uses Jini discovery/join protocols to implement its FSOOA and FMI.

In SOOA, a service provider is an object that accepts remote messages from service requestors to execute an item of work. These messages are called *service exertions* that describe service data, operations and provider's control strategy. A *task exertion* is an elementary service request, a kind of elementary remote instruction (elementary statement) executed by a single service provider or a small-scale federation. A composite exertion called a *job exertion* is defined hierarchically in terms of tasks and other jobs, a kind of network procedure executed by a large-scale federation. The executing exertion is a service-oriented program that is dynamically bound to all needed and currently available service providers on the network. This

collection of providers identified in runtime is called an *exertion federation*. This federation is also called an *exertion space*. While this sounds similar to the object-oriented paradigm, it really isn't. In the object-oriented paradigm, the object space is a program itself; here the exertion space is the *execution environment* for the exertion that is a service-oriented distributed program. This changes the programming paradigm completely. In the former case the object space is hosted by a single computer, but in the latter case the service providers are hosted by the network of computers.

The overlay network of service providers is called the *service provider grid* and an exertion federation is called a *virtual metacomputer*. The *metainstruction set* of the metacomputer consists of all operations offered by all service providers in the grid. Thus, a service-oriented program is composed of metainstructions with its own service-oriented control strategy and service context representing the metaprogram parameters. The service context describes the data that tasks and jobs work on. Exertion-oriented programs (metaprograms) can be created interactively [22] and allow for a dynamic federation to transparently coordinate their execution within the grid. Please note that these metacomputing concepts are defined differently in classical grid computing where a job is just an executing process for a submitted executable code with no federation being formed.

Each SORCER provider offers services to other peers [8] on the object-oriented overlay network. These services are exposed indirectly by methods in well-known public remote interfaces and considered as elementary (tasks) or compound (jobs) statements of the FSOOA [21]. Requestors do not need to know the exact location of a provider beforehand; they can find it dynamically by discovering Jini lookup services and then looking up a needed service implementing required service types. Service providers do not have mutual associations prior to the execution of an exertion; they come together dynamically (federate) for all nested tasks and jobs in the exertion. Specialized providers within the federation, or task peers, execute service tasks. Jobs are coordinated by a *rendezvous* or *job peer* called a *Jobber*, one of SORCER infrastructure services. However, a job can be sent to any service provider (peer). A peer that is not a Jobber type is responsible for forwarding the job to one of available rendezvous peers in the SORCER grid and returning results to the requestor.

Thus implicitly, any peer can handle any job or task. Once the job execution is complete, the federation dissolves and the providers disperse to seek other exertions to join. Also, SORCER supports a traditional approach to grid computing similar to those found in Condor [28] and Globus [27]. Here, instead of exertions being executed by services providing business logic for requested exertions, the business logic comes from the service requestor's executable programs that seek compute resources on the network.

Grid-based services in the SORCER environment include *Grider* services collaborating with Jobber services for traditional grid job submission, and *Caller* and *Methodor* services for task execution [13]. Callers execute conventional programs via a system call as described in the service context of a submitted task. Methoders download required Java code (task method) from requestors to process any submitted context accordingly with the downloaded code. In either case, the business logic comes from requestors; it is conventional executable code invoked by Callers with the standard Caller's service context or mobile Java code executed by Methoders with any service context provided by the requestor.

4 Federated Method Invocation (FMI)

Each programming language provides a specific computing abstraction. Procedural languages are abstractions of assembly languages. Object-oriented languages abstract elements in the application domain that refer to “objects” communicating via message passing as their representation in the corresponding solution space. The object-oriented distributed programming should allow us to describe the distributed problem in terms of the intrinsic unpredictable network problem instead of in terms of distributed objects that hide the notion of the network.

What intrinsic distributed abstractions are defined in SORCER? Well, *service providers* are “objects”, but they are specific objects—they are *network objects* with a *network state*, *network behavior*, and *network type(s)*. There is still a connection to distributed objects: each service provider looks like a distributed object (compute node) in that it has a network state, network behavior, and network types(s). Service providers act also as *network peers*; they are replicated and dynamically provisioned for reliability to compensate for network failures [18]. They can be found dynamically in runtime by type(s) they implement. They can federate for executing a specific network request called an *exertion* and perform hierarchically nested (component) exertions. An exertion encapsulates *service data*, *operations*, and provider’s *control strategy*. The component exertions may need to share context data of ancestor exertions, and the top-level exertion is complete only if all nested exertions are successful.

With that very concise introduction to the abstractions of exertion-based programming, let’s look in detail at how Federated Method Invocation (FMI) is structured.

4.1 Service Messaging and Exertions

In object-oriented terminology, a message is the single means of passing control to an object. If the object responds to the message, it has an operation and its implementation (method) for that message. Because object data is encapsulated and not directly accessible, a message is the only way to send data from one object to another. Each message specifies the name (identifier) of the receiving object, the name (selector) of operation to be invoked, and its parameters. In the unreliable network of objects; the receiving object might not be present or can go away at any time. Thus, we should postpone receiving object identification as late as possible. Grouping related messages per one request for the same data set makes a lot of sense due to network invocation latency and common errors in handling. These observations lead us to service-oriented messages called *exertions* that encapsulate both multiple *service signatures* that define operations and *service context* as data. Different types of exertions define own execution control strategies. Two basic exertion types are distinguished: elementary and composite exertion called *service task* and *service job* respectively. There are two ways of invoking exertions. In the first case, an `Exertion` can be invoked by calling `Exertion.exert(Transaction)`. The second way is explained in Subsection 4.6.

4.2 Service Signatures

An exertion initiates the dynamic federation of all needed service providers dynamically—as late as possible—as specified by signatures of top-level and nested exertions. Thus, FMI is defined as exerting signatures, which is essentially an indirect invocation of network methods specified by the exertion signatures for related service contexts. SORCER service providers and requestors usually communicate via FMI.

A service `Signature` is defined by:

- signature name— a custom name
 - service type— a Java interface name
 - selector of the service operation—an operation name of the service type
 - operation type— `Signature.Type` : `PROCESS` (default), `PREPROCESS`, `POSTPROCESS`
 - service access type— `Signature.Access` ; `PUSH` (default) direct binding to `Jobbers` or `Taskers`, or `DROP` using the `Spacer` service—shared exertion space
 - priority— used by exertion’s control strategy
 - execution time flag—if `true`, the execution time is returned in the service context
 - notifyees—list of email addresses to notify upon exertion completion
 - service attributes—requestor’s attributes matching provider’s registration attributes
- An exertion can comprise of a collection of `PREPROCESS` and `POSTPROCESS` signatures, but only one `PROCESS` signature. The `PROCESS` signature defines the binding provider for the exertion.

4.3 Exertion Types

A `Task` is the analog of a statement in conventional programming languages—here an elementary step of the exertion-oriented program. Thus, it is a minimal unit of structuring in exertion-oriented programming. If the provider responds to a `Task`, it has a method for the task's `PROCESS` signature. Other signatures associated with the `Task` provide for preprocessing and postprocessing by the same or its federating providers. An `APPEND` signature provides for the context received from the provider identified by this signature to be appended in runtime to the task’s currently processed service context. Appending a service context allows a requestor to use actual data in runtime not available to the requestor when a task is submitted. A `Task` is the single means of passing control to a `PROCESS` provider. Note that a task is a batch of operations that operate on the same service context—a `Task` shared execution state—and all operations of the `Task`, as defined by signatures, can be executed by the same provider or a group of federating providers coordinated by the `PROCESS` provider—the provider identified by the `PROCESS` signature of the `Exertion`.

A `Job` is the analog of a procedure in conventional programming languages—here a federated procedure is an exertion-oriented program. It is a composite of exertions (see Figure 4) that makeup the federated procedure. The following flow control exertion types define algorithmic logic of exertion-oriented programming: **`ServiceTask`**, **`ServiceJob`**, **`IfExertion`**, **`WhileExertion`**, **`ForExertion`**, `DoExertionThrowExertion`, `TryExertion`, `BreakExertion`, `ContinueExertion`. Currently implemented flow control `Exertion` types in SORCER are indicated above in bold.

4.4 Service Contexts

A service context is a data structure that describes service provider ontology along with related data. A provider ontology is controlled by provider vocabulary that describes objects and the relations between them in a provider's namespace within a specified service domain of interest. A requestor submitting an exertion to a provider has to comply with that ontology. In service context attributes and their values are used as atomic conceptual primitives, and complements are used as composite ones. A complement is an attribute sequence (path) with a value at the last position. An elementary context property consists of a context subject (main complement) and a set of context complements, and usually corresponds to a simple sentence of natural language.

A service context is a tree-like structure described conceptually in the EBNF conceptual syntax specification as follows:

1. context = [subject ":"] complement { complement }.
2. subject = element.
3. complement = element ";"
4. element = path ["=" value].
5. path = attribute { "/" attribute } [{ "<" association ">" }] [{ "/" attribute }].
6. value = object.
7. attribute = identifier.
8. relation = domain product.
9. association = domain tuple.
10. product = attribute { "|" attribute }.
11. tuple = value { "|" value }.
12. attribute = identifier.
13. domain = identifier.
14. association = identifier.
15. identifier = letter { letter | digit }.

A relation with a single attribute is called a *property* and is denoted as attribute | attribute. To illustrate the idea of context, let's consider the following example:

laboratory/name = SORCER: university=TTU;

university/department/name=CS;

university/department/room/number=20B;

university/department/room/phone/number=806-742-

university/department/room/phone/ext=237;

director <person | Mike | W | Sobolewski> /email=sobol@cs.ttu.edu;

where the relation person is defined as follows: person | firstname | initial | lastname.

A context leaf node, or *data node* is where the actual data resides. The service context—all context paths—denotes an application domain namespace, and a *context model* is its context with data nodes appended to its context paths. A context path is a hierarchical name for a data item in a leaf node. Note that Context can be represented as an XML document—what has been done in SORCER for interoperability—but the power of object Contexts comes from the fact that any Java object can be naturally used as a data node. In particular exertions themselves can be used as data nodes and then executed and controlled by providers to run complex iterative programs, e.g., nonlinear multidisciplinary optimization [14].

4.5 Service-to-Service (S2S) Computing

Tasks are usually executed by providers of the `Tasker` type (task peer). A `Job` contains a service context called *control context* that describes the control strategy for the job exertion. Dedicated service providers of the `Jobber` type (job peer also called rendezvous peer), interpret and execute a job's control context in terms of the job's nested exertions accordingly. A `Jobber` manages a shared context (shared execution state) for the job federation and provides a substitution for input context parameter mappings. A `Jobber` creates a federation of required service providers (`Taskers` and `Jobbers`) in runtime. A `SORCER` peer (`Servicer`) that is unable to execute an `Exertion` for any reason forwards the `Exertion` to any available `Servicer` matching the exertion's `PROCESS` signature and returns the resulting exertion back to its requestor.

All `SORCER` service providers are service peers as they implement the top-level `Servicer` interface. As a result, each `Servicer` can initiate a federation created in response to `Servicer.service(Exertion, Transaction)`. `Servicers` come together to form a federation participating in execution of the same exertion. When the exertion is complete, `Servicers` leave the federation and seek a new exertion to join. Note that the same exertion can form a different federation for each execution due to the dynamic nature of looking up `Servicers` by their implemented custom interfaces. Despite the fact that every `Servicer` can accept any exertion, `Servicers` have well defined roles in `SORCER` S2S exertion-oriented programming:

- a) `Taskers` – process service tasks
- b) `Jobbers` – process service jobs
- c) `Contexters` – provide service contexts for `APPEND` Signatures
- d) `FileStorers` – provide access to federated file system providers [1, 23]
- e) `Cataloger` – `Servicer` registries
- f) `Persisters` – persist service contexts, tasks, and jobs to be reused for interactive exertion-based programming
- g) `Spacers` – manage exertion spaces shared across `Servicers` for space-based computing [7]
- h) `Relayers` – gateway providers, transform exertions to native representation, for example integration with Web services and JXTA
- i) `Authenticators`, `Authorizers`, `Policers`, `KeyStorers` – provide support for service-oriented security
- j) `Auditors`, `Reporters`, `Loggers` – support for accountability, reporting and logging
- k) `Griders`, `Callers`, `Methoders` – support conventional grid computing
- l) `Generic ServiceTasker` and `ServiceJobber` implementations are used to configure domain specific providers via dependency injection—configuration files for smart proxying and inserting business objects called `SORCER` service beans.

4.6 FMI Triple Command Pattern

Polymorphism lets us encapsulate a request—an exertion—then establish the signature of operation to call and vary the effect of calling the underlying operation by varying its implementation. The Command design pattern [10] establishes an operation signature as an interface and defines various implementations of the interface. In FMI, the following three operations are defined:

1. `Exertion.exert(Transaction):Exertion`—join the federation;
2. `Servicer.service(Exertion, Transaction):Exertion`—request a service in the federation from the top-level `Servicer` obtained by the receiver;
3. `Exerter.exert(Exertion, Transaction):Exertion`—execute the argument exertion by the target provider in the federation.

The above *Triple Command* pattern defines various implementations of these interfaces: `Exertion`, `Servicer`, and `Exerter`. This approach allows for the P2P environment via the `Servicer` interface, extensive modularization of `Exertions` and `Exerters`, and extensibility from the triple design pattern so requestors can submit any service-oriented programs (exertions) they want with or without transactional semantics. FMI triple Command Pattern is used as follows:

1. An exertion can be invoked by calling `Exertion.exert(Transaction)`. The `Exertion.exert` operation implemented in `ServiceExertion` uses `ServicerAccessor` to locate in runtime the provider matching the exertion's `PROCESS` signature.
2. If the matching provider is found, then on its access proxy (that can also be a smart proxy) the `Servicer.service(Exertion, Transaction)` method is invoked.
3. When the requestor is authenticated and authorized by the provider to invoke the method defined by the exertion's `PROCESS` signature, then the provider calls its own `exert` operation: `Exerter.exert(Exertion, Transaction)`.
4. `Exerter.exert` method calls `exert` either of `ServiceTasker` or `ServiceJobber` (depending on the type of the exertion: either `Task` or `Job`) that by reflection calls the method specified in the `PROCES` signature (interface and selector) of the exertion. All application domain methods of any application interface (custom `Tasker` interfaces) have the same signature: a single `Context` type parameter and a `Context` type return vale. Thus a custom interface looks like an RMI interface with the above simplification on the common signature for all interface methods.

In the FMI approach, a requestor can create any `Exertion`, composed from any hierarchically nested `Exertions`, with any service provider supplied ontology. The provider's service context ontology, object proxies and registration object attributes are network-centric; all of them are part of the provider's registration so they can be accessed via `Cataloger` or lookup services by any requestor on the network, e.g., service browsers [11] or custom service UI user agents [29] providing interactive exertion-oriented programming. In `SORCER`, using these zero-install service UIs, the user can define data for downloaded ontology and create a task/job to be executed on the virtual metacomputer.

Individual Providers, in particular Taskers and Jobbers, implement their own `exert(Exertion, Transaction)` methods according to their service semantics and control strategy; in SORCER implemented by `ServiceTasker` and `ServiceJobber` respectively. SORCER specific domain providers either subclass `ServiceTasker` or `ServiceJobber`, or by dependency injection (using Jini configuration methodology) configure either one with one of 12 proxying methods developed in SORCER. In general, many different types of taskers and jobbers can be used in SORCER at the same time (currently one `ServiceTasker` and one `ServiceJobber` implementation exists) and exertions via their signatures will make appropriate runtime choices as to what virtual metacomputer to run.

Invoking an exertion, let's say `ext`, is similar to invoking an executable program `ext.exe` at the command prompt. If we use the Tenex C shell (`tcsh`), invoking the program is equivalent to: `tcsh ext.exe`, i.e., passing the executable `ext.exe` to `tcsh`. Similarly, to invoke a metaprogram using FMI, in this case the exertion `ext`, we call `ext.exert(null)` if no transactional semantics is required. Thus, the exertion is the metaprogram and the network shell at the same time, which might first come as a surprise, but close evaluation of this fact shows it to be consistent with the meaning of object-oriented federated programming. Here, the *virtual metacomputer* is a federation that does not exist when the exertion is created. Thus, the notion of the *virtual metacomputer* is enclosed in the exertion exemplified by FMI.

The observation concluding that the exertion is the metaprogram and the network shell at the same time brings us back to the distribution transparency issue discussed in Section 2. It might appear that `Exertion` objects are network wrappers as they hide network intrinsic unpredictable behavior. However, `Exertions` are not distributed objects, as do not implement any remote interfaces; they are local objects. `Servicers` are distributed objects and there are many types of `Servicers` addressing different aspects of networking. The network intrinsic unpredictable network behavior is addressed by the SORCER object-oriented distributed infrastructure: `Taskers`, `Jobbers`, `Catalogers`, `Spacers`, `FileStorers`, `Authenticators`, `Authorizers`, `Policers`, etc. The `Servicer`-based infrastructure facilitates exertion-oriented programming and concurrent metaprogram execution using the presented FMI and allows for constructing reliable object oriented distributed systems from unreliable distribute components - `Servicers`.

5 Conclusions

A distributed system is not just a collection of distributed objects—it's the network of dynamic objects. From an object-oriented point of view, the network of dynamic objects is the *problem domain* of object-oriented distributed programming that requires relevant abstractions in the *solution space*. The exertion-based programming introduces the new RPC abstraction with *service providers* and *exertions* instead of object-oriented conventional *objects* and *messages*. Service providers can register proxies, including smart proxies, via dependency injection using twelve methods investigated in SORCER. Executing a top-level exertion means forming a dynamic

federation of currently available providers in the network that collaboratively processes service contexts of all nested exertions. Services are invoked by passing exertions on to providers indirectly via service object proxies that in fact are access proxies allowing for service providers to enforce a security policy on access to services. When a permission is granted, then the operation defined by an exertion's PROCESS signature is invoked by reflection. FMI allows for the P2P environment via the *Service* interface, extensive modularization of *Exertions* and *Exerters*, and extensibility from the triple command design pattern. The presented FMI has been successfully deployed and tested in multiple concurrent engineering and large-scale distributed applications.

References

1. Berger M., and Sobolewski M., *SILENUS – A Federated Service-oriented Approach to Distributed File Systems*, In Next Generation Concurrent Engineering, ISPE/Omnipress, pp. 89-96 (2005).
2. Birrell A. D. & Nelson B. J., *Implementing Remote Procedure Calls*, XEROX CSL-83-7, October 1983.
3. Edwards W. K., *Core Jini*, 2nd ed., Prentice Hall (2000).
4. *Fallacies of Distributed Computing*, Accessed on: July 15, 2007. Available at: http://en.wikipedia.org/wiki/Fallacies_of_Distributed_Computing
5. *FIPER: Federated Intelligent Product Environment*, Available at: <http://sorcer.cs.ttu.edu/fiper/fiper.html> Accessed on: July 15, 2007.
6. Foster I., Kesselman C., Tuecke S., *The Anatomy of the J. Supercomputer Applications*, 15(3) (2001).
7. Freeman E., Hupfer S., & Arnold K. *JavaSpaces™ Principles, Patterns, and Practice*, Addison-Wesley, ISBN: 0-201-30955-6 (1999).
8. Goel S., Shashishekara Talya S. S., Sobolewski M., *Service-based P2P overlay network for collaborative problem solving*, Decision Support Systems, Volume 43, Issue 2, March 2007, pp. 547-568 (2007).
9. Goel S, Talya S., and Sobolewski M., *Preliminary Design Using Distributed Service-based Computing*, Proceeding of the 12th Conference on Concurrent Engineering: Research and Applications, ISPE, Inc., pp. 113-120 (2005).
10. Grand M., *Patterns in Java*, Volume 1, Wiley, ISBN: 0-471-25841-5 (1999).
11. *Inca X™ Service Browser for Jini Technology*, Available at: <http://www.incax.com/index.htm?http://www.incax.com/service-browser.htm>
12. Jini architecture specification, Version 2.1. Available at: <http://www.sun.com/software/jini/spec/jini1.2html/jini-title.html> Accessed on: March 15, 2007 (2001)
13. Khurana V., Berger M., Sobolewski M., *A Federated Grid Env. with Replication Services*, In Next Generation Concurrent Engineering, ISPE/Omnipress (2005).
14. Kolonay R. M., Sobolewski M., Tappeta R., Paradis M., Burton S. 2002, *Network-Centric MAO Environment*, The Society for Modeling and Simulation International, Westrn Multiconference, San Antonio, TX (2002)
15. Lapinski M., Sobolewski M., *Managing Notifications in a Federated S2S Environment*, International Journal of Concurrent Engineering: Research & Applications, Vol. 11, pp. 17-25 (2003).
16. McGovern J., Tyagi S., Stevens M. E., Mathew S., *Java Web Services Architecture*, Morgan Kaufmann (2003).

17. Pitt E., McNiff K., *java.rmi: The Remote Method Invocation Guide*, Addison-Wesley Professional (2001).
18. Project Rio, A Dynamic Service Architecture for Distributed Applications. Available at: <https://rio.dev.java.net/>. Accessed on: March 15, 2007.
19. Röhl P. J., Kolonay R. M., Irani R. K., Sobolewski M., Kao K. *A Federated Intelligent Product Environment*, AIAA-2000-4902, 8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, Long Beach, CA, September 6-8 (2000).
20. Ruh W. A., Herron T., Klinker P., *IOP Complete: Understanding CORBA and Middleware Interoperability*, Addison-Wesley (1999).
21. Sobolewski M., *Federated P2P services in CE Environments*, Advances in Concurrent Engineering, A. A. Balkema Publishers, 2002, pp. 13-22 (2002).
22. Sobolewski M., Kolonay R., *Federated Grid Computing with Interactive Service-oriented Programming*, International Journal of Concurrent Engineering: Research & Applications, Vol. 14, No 1., pp. 55-66 (2006).
23. Sobolewski M., Soorianarayanan S., Malladi-Venkata R-K. 2003, *Service-Oriented File Sharing, Proceedings of the IASTED Intl., Conference on Communications, Internet, and Information technology*, pp. 633-639, ACTA Press (2003).
24. Soorianarayanan S., Sobolewski, M., *Monitoring Federated Services in CE*, Concurrent Engineering: The Worldwide Engineering Grid, Tsinghua Press and Springer Verlag, pp. 89-95 (2004).
25. *SORCER Research Group*, Available at: <http://sorcer.cs.ttu.edu/>
26. *SORCER Research Topics*, Available at: <http://sorcer.cs.ttu.edu/theses/>
27. Sotomayor B., Childers L., *Globus® Toolkit 4: Programming Java Services*, Morgan Kaufmann (2005).
28. Thain D., Tannenbaum T., Livny M. *Condor and the Grid*, In Fran Berman, Anthony J. G. Hey, and Geoffrey Fox, editors, *Grid Computing: Making The Global Infrastructure a Reality*. John Wiley (2003).
29. *The Service UI Project*, Available at: <http://www.artima.com/jini/serviceui/index.html>. Accessed on: July 15, 2007.
30. Waldo J., *The End of Protocols*, Available at: <http://java.sun.com/developer/technicalArticles/jini/protocols.html>. Accessed on: March 15, 2007.