

A Design Pattern for the Transfer of Running Applications between Devices

Paul Austrem
University of Bergen
Fosswinckels gate 6
5007 Bergen, Norway
paul.austrem@infomedia.uib.no

Abstract. Mobile information systems are growing in acceptance; in order for the vision of true mobility to be realized users must be able to seamlessly move running applications between devices in an *ad-hoc* manner. The task of implementing mechanisms to ensure that a running application is successfully moved from one device to another can be considered a generic task, wherein the same fundamental design can be reused. This work tenders a domain specific design pre-pattern that resolves the problems associated with transferring a running application from one device to another whilst maintaining state and tailoring to capabilities. It is meant to aid in the domain of mobile workers, and can be applied as a container between the OS and the applications or as a value-added service by ASPs. The solution adds complexity and imposes conventions on the extendibility of a system, but allows users to maintain state so they can seamlessly move their work between devices.

1 Overview

More and more workers are befitting the term mobile knowledge workers. These are workers who in order to perform their job need to have timely access to information and applications, and may be performing their work in places that have not been determined upfront. Even more so mobile workers may need to suddenly transfer their ongoing work from one device to another. These are functional challenges that are rooted in the evolvement of mobile computing and need to be resolved if the promise of true mobility is to come to fruition. This will be a recurring functional challenge, as mobile workers across domains and industries will wish to switch execution of their application from one device to another in an *ad-hoc* manner as the circumstances for them completing their work changes. For instance a journalist may need to leave her desk, however she wishes to continue her current work so she transfers the application she is running from her desktop computer to her mobile device. Since she is working with text editing she is very keen on maintaining the state, the operational history, of the application on the new device, thus maintaining access to for example “undo” operations as they were performed on the desktop computer. This pattern ensures that the application is transferred maintaining this state.

The IT business analysis company Gartner has presented such functionality as being the next big evolution in middleware and mobile information system solutions [1]. This work presents a domain specific design pre-pattern solution to such a problem of transferring a running process between heterogeneous devices. The design pattern is intended for use in scenarios

involving mobile workers. It will be presented using a subset of the design pattern format used by the GoF [2] and accompanied by a class diagram, an activity diagram and a sequence diagram.

“PROCESS ON THE GO”

2 Intent

Mobile workers conduct their business in changing environments and with changing resources. They need to be able to quickly move their work from one device to another in a seamless manner that does not interrupt their flow. There should be transparency between devices; there should be no “scars” indicating that the work has been transferred between multiple devices during a work session.

Therefore:

Provide a means to handle the transfer of a running application from one device to another whilst maintaining application state and adapting to the target device’s performance profile, thereby enabling seamless mobility.

“PROCESS ON THE GO” handles two separate tasks; the transfer of state between devices and the tailored reconstruction of the application on the target device. State is represented as a structure containing all the operations performed by a user during an uninterrupted user session. An interruption is the termination of the session through the application being closed.

3 Motivation

A mobile worker, Jill, is preparing a sales presentation to be held in a distant city. Unfortunately before she can finish it she must leave for the airport. She transfers the application to a mobile device knowing that she will spend the next 30 minutes travelling as a passenger, thereby being able to work on the presentation. When she reaches the airport she transfers the application to her laptop so she can continue work during travel. The pattern essentially makes it completely transparent in terms of the state and representation of the application whether it has been transferred between devices. The pattern allows for the state of an application to be tailored to the target device and context it will be used in. If certain functions are not available due to computational limitations on the target device then these can be disabled in the state during the transfer. Alternatively state could be disabled due to security reasons, for instance undoing financial transactions is disabled on a mobile device being used in a public area but allowed on a stationary computer. To attain this capability one can “tag” elements of the State structure to indicate they are disallowed on the current device.

4 Applicability

The pattern can be applied to systems where there is a distinct chance that workers will use the system in an *ad-hoc*, mobile and unpredictable manner. Solutions that are used by workers

that move around a lot during their workday, or do not upfront know when or where they will be conducting certain business tasks could benefit from this design pattern. Currently primitive solutions prevail; e.g. users could store their work on a server, or simply just copy files from one device to another to continue work. However such primitive solutions, although simple, do not maintain the application's state in terms of operation/user action history. The transfer of the work task from one device to another becomes stateless and manual. "PROCESS ON THE GO" resolves this issue. Use the pattern when:

- Users would benefit from retaining the internal state of their applications across devices
- You wish to establish a framework or create a middleware solution wherein all devices follow interfaces that enable their interoperability.

The pattern is useful in applications where operational history is important, such as the undo operations in a text document or maintaining the contents of the clipboard. Additionally the pattern allows one to tailor the reconstruction of a process based on a mix between the profile and capabilities of the target device. For instance if the target device is a mobile phone, and the user profile of the mobile phone is set to "speech only", then this information can be used by the server to tailor the representation of the reconstructed application on the target device.

Following Jill's travel, she has now arrived at her destination and has hired a car as she will be spending some time travelling between various meetings. Unfortunately she didn't quite manage to finish her presentation whilst travelling, and her laptop's battery is almost drained. Therefore, she sets her mobile phone to "speech only" mode and transfers the application to it. The server creates a profile of the target device (in this case the mobile phone) and tailors the representation on the target device. For instance due to the limited screen size and resolution of the mobile phone Jill will not be able to see the presentation slides on the screen, only pure text. Furthermore she will not have access to the operational history that involves adding/editing/deleting elements that cannot be represented on the current device (e.g. images, animations, etc.). This to hinder that she inadvertently makes changes that are not visible to her. Whilst driving she can now use voice commands to work with the presentation. Using voice commands she narrates her presenter's notes to specific slides and saves the file.

5 Structure

The *client* class (fig 1) represents the source device that is the device on which the application is currently running. It implements the *IDeviceProfile* interface thereby being applicable for use in an application transfer action. The *client* class has the responsibility of maintaining its own internal state up until the point where transfer commences. *Client* in this context is different from the traditional client/server roles of for instance the world wide web. *Client* merely denotes the source device, the device from which the application will be transferred.

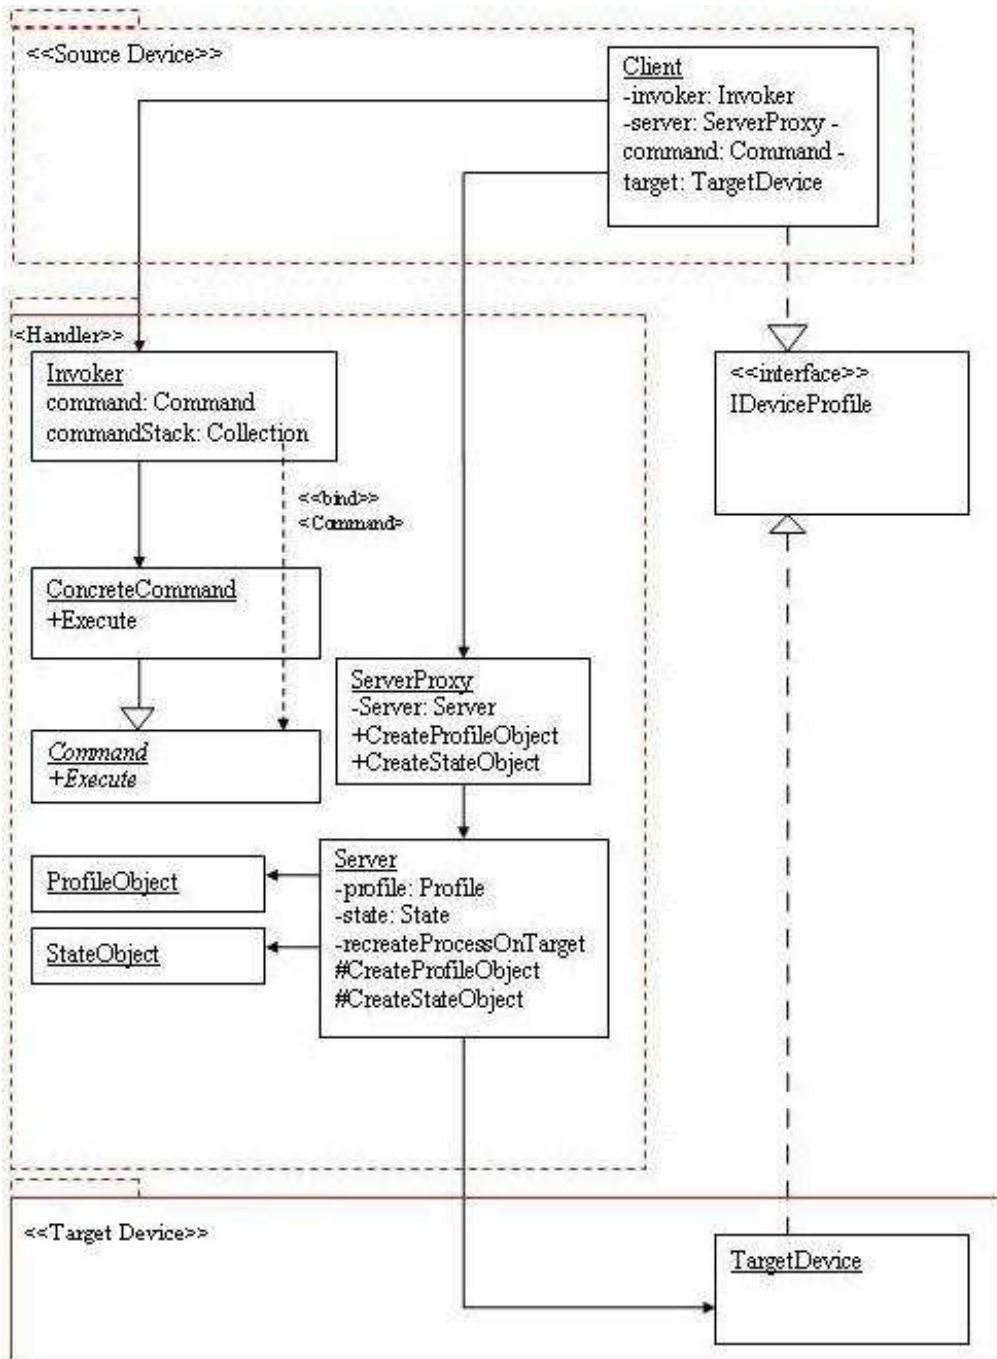


Fig 1: Class diagram of the "PROCESS ON THE GO" pattern

The classes *Invoker*, *ConcreteCommand* and *Command* are associated with the *client* class and are the classes utilized to maintain the operational history of an application, wherein the *Invoker* maintains a collection, known as the "Command Stack" which is the representation of State, with a "Last In First Out" structure of all operations / user actions.

ServerProxy is used as a front-end for the *Handler* package and used by the *Client* to provide the required objects to perform the transfer. This *ServerProxy* also enforces the decoupling and opaqueness of the client and handler, since the client only communicates with a proxy

object it has no real knowledge of where the server actually resides (is it a dedicated handler, a handler on the client itself, or on the target device?). The *ServerProxy* does not perform any operations itself, it merely decomposes and delegates the tasks on to the *Server* class.

The *Server* class performs the required processing in order to transfer the process to the target device. It uses the two structural classes *ProfileObject* and *StateObject* to maintain the data required to reconstruct the process with state on the target device. The *ProfileObject* structural class is essentially a message format, defining the structure of a profile message. It may describe the QoS characteristics of the target device, for instance the CPU power, multi-threading capabilities, screen resolution etc. The object is used by the *recreateProcessOnTarget* method. Since this object follows a predetermined format it establishes a shared ontology between the devices and the server thus allowing the server to understand the capabilities of devices and tailor the application recreation accordingly.

The *StateObject* is a structure containing a *CommandStack*, a list of all user operations performed during the current session. The Journaling pattern [2] may be used to improve performance in terms of adding new Commands to the *CommandStack*. Especially true if the *CommandStack* is stored as a flat file involving disk operations.

Finally the *TargetDevice* class represents the target device on which the application will be recreated. It implements the *IDeviceProfile* which allows it to be "profiled" by the *Server*. This runtime profiling, as opposed to the server maintaining a database of device profiles strengthens the decoupling between the server and the target devices, thus the server can transfer between devices that it originally did not know existed.

Furthermore, the pattern is a composite pattern and utilizes several other patterns in order to attain its objective. We can identify the applicability of the "COMMAND" [3] pattern as it maintains the operations history; the user actions that have been performed in the application. The "COMMAND" pattern's "Caretaker" object would reside on the handler side (figure 1). The rationale behind this is that if there is in fact a physical separation between the source device and the handler, then in case of the source device experiencing a terminal exception when moving the process the object state would be stored separately. Hence the state could be recreated when the application is reinitiated on the source device, or the transfer of the process could continue to the target device.

The other task we are concerned with here is the use of three patterns to enable the creation of the *ProfileObject*. As we can see from figure 2 there are a set of design requirements that support the use of the "INVOKER" [4] pattern, the "INTERFACE DESCRIPTION" [4] pattern and the "OBJECT ID" [4] pattern. Firstly the use of "OBJECT ID" is warranted because in a mobile work environment a user may have the possibility to move her work process to several different mobile devices. For instance the process could be transferred from a stationary computer to either a laptop computer or an ultra-portable PC or a PDA. The pattern ensures that the "server" invokes the retrieval of the profile from the correct remote target device. This leads onto the use of the "INTERFACE DESCRIPTION" pattern. The "poll target device" activity encourages the use of this pattern, since in order for various devices to be able to poll each other's capabilities it is a pre-requisite that they share a pre-agreed set of methods that can be used for this purpose. The "INTERFACE DESCRIPTION" pattern supports this as both the client device and the target device will be forced to adhere to the method signatures defined in a shared interface. Finally the "INVOKER" pattern will be used to enable the actual communication between the remote objects; a pre-requisite in this case is the use of "OBJECT

ID” to ensure the client device has the required ID of the target device. The client device acts the role of the “Requestor” in the “INVOKER” pattern, whereas the server (ref fig 1) acts the role of the “Invoker”. Thus when the client delegates the task of retrieving the target device profile to the “ServerProxy” it passes in a TargetDevice object which contains the signature/objectID of the target device to the server through the server proxy object. The “ServerProxy” class is used to decouple the client from the actual server, as mentioned; depending on the environment in which the client device is working there may, or may not be access to a physically separate server. However, the client application should work without any explicit knowledge of this. Therefore the proxy class is used to enforce this opaqueness.

After the server has asynchronously created the two objects “profile” and “state”, which are essentially just structures, it will initiate the operation “recreateProcessOnTarget” in which it will recreate the application on the target device based on information from the “profile” and “state” object.

6 Collaborations

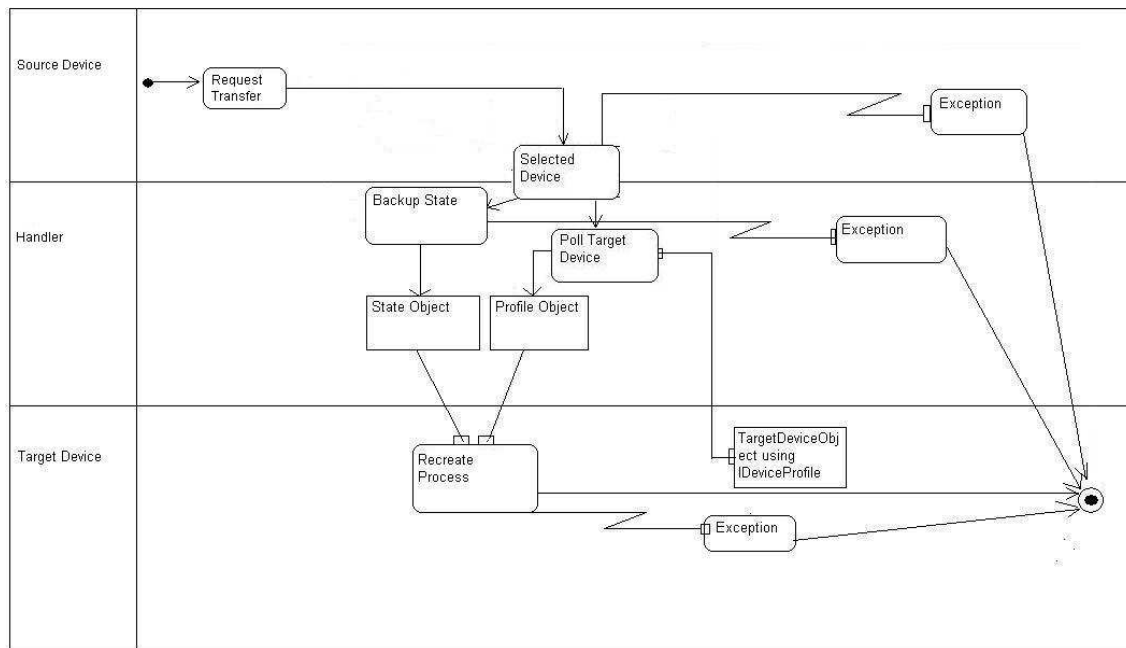


Fig 2: An activity diagram showing a generic approach to moving a running process from one device to another

Firstly we use activity partitions to create swimlanes in which each of the partaking devices is positioned. In this approach it is feasible to introduce three swimlanes denoting the client, the server and the target device. The *Source Device - Handler* division is purely logical, since both the *Source Device* and the *Handler* could potentially exist on the same device, or the handler could even exist on the target device, thus the *Handler* is not necessarily a separate physical entity or disk; it is a logical description that separates disparate tasks associated with transferring an application between devices. The description of *Handler* is a role, a part played for a short duration before it is passed on. As soon as the application has been recreated on the target device the command stack is deleted from the *Handler* and the application is closed on the *Source Device*. The *Target Device* then becomes the new *Source Device*. If the recreation had failed however, the *Handler* notifies the original *Source Device*

of this and “rolls-back” any operations performed on the *Target Device*. This approach is simplistic and disallows one to transfer an application to multiple devices simultaneously. However, the simplicity also resolves issues that stem from multiple instances of the same application at the same time and also succeeds challenges with merging different command stacks from different devices. Thus there is never more than one instance of the command stack at any given time, and all contents of the stack are preserved. If certain operations in the stack cannot be performed due to device limitations or business rules then they are “tagged” in the stack, but never removed. Thus the Command Stack maintains its consistency across devices regardless of their profile and capabilities.

The *Handler* sustains a backup of the state of the process before it is attempted moved. This backup should preferably be stored on a persistent storage device, e.g. a memory card in the mobile device or a hard-drive in a laptop or even on a separate server; however in-memory storage is also acceptable if no other viable options are available. The point is that this data should be logically separated from the process working on it, which would be the source device application. Therefore the backup state and state object reside at the *Handler* level.

Sequence Diagram

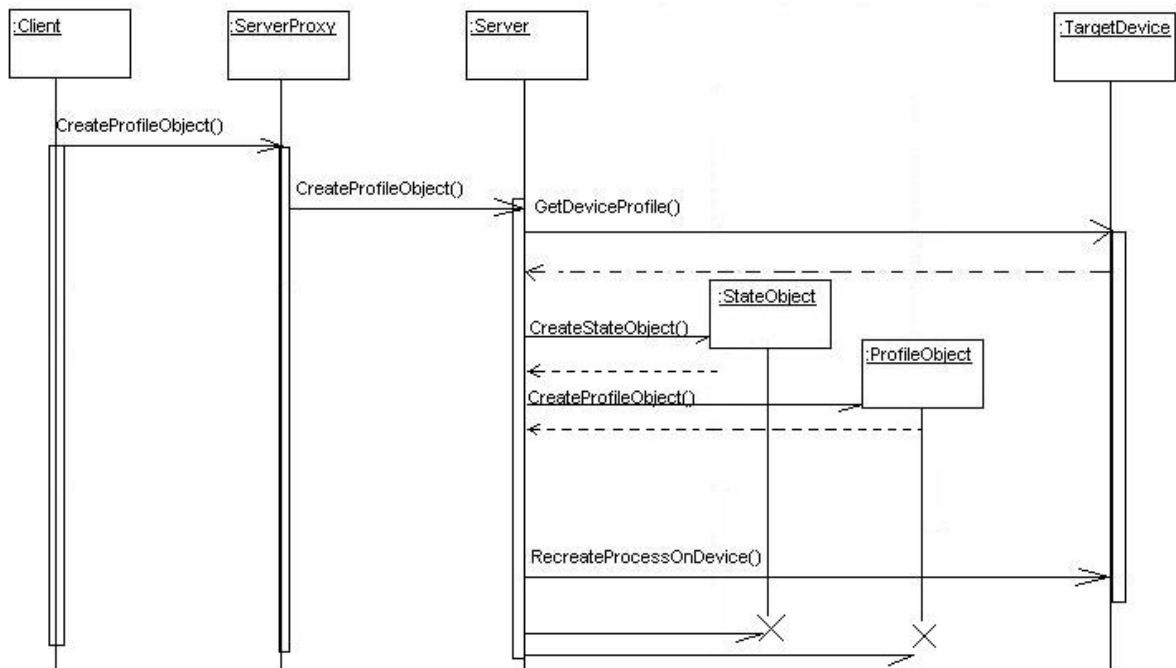


Fig 3: Sequence Diagram of the “PROCESS ON THE GO” pattern

The client initiates the process by contacting the handler through *ServerProxy*, using two separate requests to pass the required data to construct the *Profile* object and the *State* object. The *ServerProxy* forwards the requests to the actual server that creates a *Profile* object and a *State* object. When both objects have been successfully created the *Server* calls the *CreateProcessOnDevice* method, which is defined in the *IDeviceProfile* interface implemented by the *TargetDevice*.

Considerations

One issue that needs to be resolved is how the command stack is handled by devices with limited capabilities. If a mobile device can only perform a subset of the operations performed on a desktop computer there will be a need to handle this capability mismatch. The server can use the *ProfileObject* to “tag” operations in the Command Stack (*StateObject*) that cannot be performed due to the limitations of the target device. This way when recreating the state the target device will skip these tagged operations; they will still remain in the stack but be unavailable as long as the user is working on the limited device. Hence, if an application is transferred from a laptop to a limited device and back again no operations will be lost, the stack will be intact because for each transfer the stack will be “re-tagged”. However there are exceptions to this principle. For instance considering proprietary software developed for multiple devices using a MVC approach, wherein all applications regardless of device share the same “Controller”, or functional core. Thus, the “View” is adapted based on the profile of the device, but the functionality of the “Controller” is maintained. This would allow the user to perform all operations on the stack, for example to “undo” operations on a graphical entity in the application although the same operations are not directly supported through the user interface of the application.

Another issue that should be considered is security, for instance certain functions should not be available, or be “undoable” when the application is executed on a mobile device. We could imagine a mobile worker undoing financial transfers; business rules dictate that such operations are only permitted on a user’s stationary computer. However the user may still be allowed to perform work in other parts of the application, for example fill out an electronic form etc. This could be solved by the handler tagging certain operations in the stack as unavailable, before transferring and recreating the application on the target device.

Jill has finally arrived at her destination, and is now almost ready to give her presentation. She has plugged in the AC adapter for her laptop and is recharging it. Since she will be running her presentation from her laptop she needs to transfer the presentation back from her mobile phone to the laptop. This time the profile built by the server indicates that the target device is highly capable both in terms of power and presentation alternatives (screen, sound). Thus the server restores the presentation on the target device (laptop) with all functions and history available allowing Jill to add her finishing touches before presenting.

7 Implementation Considerations

Ideally the pattern should not affect the existing architecture; therefore it shouldn’t be a native function of any applications. This paper *describes* two manners in which it could be implemented; both have their advantages and drawbacks. The first approach involves an application service provider (ASP) wherein a mobile worker may license a software product across multiple devices. The ASP stores the operational history, the state, of the application – thus the pattern is implemented with a dedicated physical handler, and the mobile worker accesses the application through a thin-client, the worker’s device. The strength of this is naturally that this does not require any adaptations from the mobile worker using the product as it is all handled by the ASP. Whenever the mobile worker performs an operation on the

ASP web-based application a network call, e.g. a HTTP call if it uses browser-based access, is submitted containing information about the operation. The calls can then be stored with timestamps and ID of the user in a database. The operational history can then easily be restored whenever the mobile worker switches between devices. A drawback is that one would need an internet connection in order to utilize the seamless mobility.

The other alternative would be to implement the pattern as a middleware solution installed on all devices that are used by the mobile worker. It works locally registering the operations performed by the user, and adding them to the command stack. Every running application is added to this transparent container, as mentioned a key aim is to make the pattern transparent to any applications contained in it. Thus, no changes are implemented on the platform or application.

8 Consequences

The following general advantages are provided by the “PROCESS ON THE GO” pattern:

- *Workers that use “complex” software where maintaining state is important can become truly mobile.* Complex business processes that may involve multiple applications can be transferred seamlessly between devices allowing workers to perform their work anytime anywhere.
- *Convention over configuration.* The core-system does not have to be configured to accommodate new devices as long as they adhere to the IDeviceProfile interface. This makes it easy to extend.

The pattern will affect the non-functional characteristics of the system in the following manner:

- *Reduced performance:* Performance will likely suffer as recording all operations and user actions will undoubtedly require additional time and resources, in addition the actual process of transferring a process from one device to another will also require resources, thus the more complex the process is, the more performance will suffer.
- *Increased ubiquity:* A system implementing this pattern will become more ubiquitous shifting the focus away from technical limitations towards user mobility.
- *Increased complexity:* The actual implementation complexity will likely be increased as developers must write the *Server* logic profiling operations.
- *Reduced flexibility:* Although this pattern is only one way of handling the transferral of applications from one device to another, the “convention over configuration” axiom used reduces the overall flexibility as all devices must implement a certain set of methods as defined in the IDeviceProfile interface, and it is only these methods that will be used in the transfer process.
- *Increased extendibility:* A converse effect of the causes of reduced flexibility is an increase in extendibility. It is easy to add new devices as all requirements are specified and determined up front through the use of the interface. All devices are essentially autonomous and have no deep knowledge of, nor interest in, the other devices. This is ensured through the decoupling provided by the server.
- *Command stack:* The command stack may become problematic if it grows too large thus not running well on devices with limited resources.

References

1. Gartner: "Gartner Says Business Application Vendors Face Challenge to Move to 'The Process of Me'". Gartner Press Releases, 23. May 2006. Gartner Symposium/ITxpo. Online resource at <http://www.gartner.com/it/page.jsp?id=492897>. (2006)
2. PerlDesignPatterns TinyWiki: "Journaling Pattern". Online resource at <http://perldesignpatterns.com/?JournalingPattern>
3. Gamma, E., Helm, R., Johnson, R. and Vlissides, J.: "Design Patterns – Elements of Reusable Object-Oriented Software" Addison Wesley (1995)
4. Völter, M., Kircher, M. and Zdun, U.: "Remoting Patterns: Foundations of Enterprise, Internet and Realtime Distributed Object Middleware". Wiley Series in Software Design Patterns (2004)
5. OMG: "UML Superstructure Specification 2.1.1". Online resource at <http://www.omg.org/technology/documents/formal/uml.htm>. (2004)