

# A Model of Dynamic Binding in .NET

Alex Buckley  
Imperial College London  
a.buckley@imperial.ac.uk

May 2005

## Abstract

Millions of programmers use ECMA CLI-compliant languages like VB.NET and C#. The resulting bytecode can be executed on several CLI implementations, such as those from Microsoft and the open-source Mono organisation. While assemblies are the standard unit of deployment, no standard exists for the process of finding and loading assemblies at run-time. The process is typically complex, and varies between CLI implementations. Unlike other linking stages, such as verification, it is visible to programmers and can be a source of confusion.

We offer a framework that describes how assemblies are resolved, loaded and used in CLI implementations. We strive for implementation-independence and note how implementations from different organisations vary in behaviour. We describe the reflection features available for dynamic loading, and give C# examples that exercise the features modelled in the framework.

## 1 Introduction & Motivation

Traditional language mechanisms for modular development - packages in Ada, modules in Modula-2, namespaces and classes in C++ - have no role at run-time. A compiler typically employs a static linker to emit a monolithic executable file, so the compilation environment automatically becomes the entire execution environment. Few (or no) dynamic checks are needed to resolve external dependencies. In contrast, the basic unit of development in Java and C# - the class - maintains its discrete identity throughout compilation and execution. A Java Virtual Machine or Microsoft's .NET Common Language Runtime (CLR) can start with the bytecode for just one class, then lazily load and link classes from the execution environment as necessary for continued execution.

Class loading tends to be highly configurable, unlike later linking stages such as verification. UNIX offered the `dlopen` C-language interface and most OO languages offer an API for dynamic class loading in their reflection libraries. Java has the familiar `CLASSPATH` mechanism for identifying class locations, and custom classloaders can be installed into the JVM.

Matters are complicated in a CLI implementation [ECM02] because classes are not deployed as standalone units. Instead, classes are encapsulated inside assemblies. An assembly enumerates classes it provides and also the names of other assemblies whose classes it uses. Assembly resolution consists of converting a bytecode's reference to an assembly name into a physical location where a suitable assembly exists. Because an assembly's identity incorporates version and security information, resolving an assembly is more complex than (and indeed, a pre-requisite to) finding a class inside an assembly.<sup>1</sup>

Different CLI implementations have different rules for assembly resolution. Also, the process of loading an assembly from a given location is implementation-specific. We use the term *binding* for the combined resolution and loading process.

---

<sup>1</sup>We do not consider, in this paper, the resolution of classes in an assembly or of members in a class.

Microsoft’s CLI implementation, the Common Language Runtime (CLR), [MG00] provides a user-configurable, network-aware system for binding called “Fusion”. For resolution, it supports a hierarchy of policies that can modify the requested version of an assembly. This allows security and performance patches to be used without rebuilding an assembly’s manifest (akin to recompiling source code). It can also resolve references to assemblies compiled for other platforms, such as the .NET Compact Framework. For loading, it supports downloading of code from remote machines and, as a last resort, on-demand installation where the user is asked to provide an assembly.

Fusion’s behaviour is typically explained in verbose official technical documentation. Recently, “blogs” written by Microsoft employees [Coo05, Shi05, Zha05, Far05, Pra05] have explained areas of poorly documented behaviour in the current CLR release (v1.1), and given detailed information about the next CLR release (v2.0). Programmer understanding is significantly enhanced by this new channel, but there is no single place where dynamic loading is explained in full detail from ‘top to bottom’. One must piece together information from around the Internet in order to explain a program’s exact assembly and classloading behaviour.

An alternative CLI implementation is Mono [dI05]. Its functionality is a subset of the CLR’s, including for assembly binding, so documentation is shorter and easier to understand. However, its binding process is subtly different from Fusion’s. Other CLI implementations, such as Microsoft’s .NET Compact Framework for mobile devices, also exhibit different behaviour from Fusion.

We wish to unify the rules that govern assembly binding in CLI implementations. We present a model that describes, at one level, how assemblies are bound (*i.e.* resolved and loaded), and at another level, how loaded assemblies are used when evaluating bytecode instructions. Assembly binding is interleaved with bytecode evaluation as in all current CLI implementations.<sup>2</sup> The model is parameterised by resolution and loading policies, so we specialise it for the Microsoft CLR v1.1 and Mono v1.1.

## 2 The Assembly Model

### 2.1 Assembly structure

In the COM and Java environments, a file that contains code has only one identifier: its filename. “DLL hell” [EJS02] arises because multiple DLL files, each containing different code, share the same filename and are placed in a shared location on disk. An application’s dependency is resolved to a filename, but there is no guarantee that the DLL file with that filename is what the application was tested against. Java applications face a similar problem, even without a standard location in the filesystem for classes.

In contrast, the CLI specification [ECM02] gives an assembly a logical identity quite different from its filename. We call this identity an *assembly name*, and reflection APIs in CLI-compliant languages typically make it a first-class value. It contains a display name, a version number (consisting of major, minor, revision and build numbers), a cultural identifier (for internationalisation) and a public key. It is convenient to just consider the presence of a security value in an assembly name, rather than the public key per se.

*AssemblyName*  $\alpha$  :  $AN =$

*DisplayName* :  $id$ , *Version* :  $int \times int \times int \times int$ , *Culture* :  $id$ , *Security* :  $id$ , *Retargetable* :  $bool$

Binding maps an assembly name to an assembly definition. All elements in an assembly name are potentially used during binding, *e.g.* if the culture is present, it can be used to choose a directory on disk where an assembly definition might be found. The security value plays the most important role because it determines whether an assembly name is a *strong name*. A non-null security value indicates that the assembly has been signed by a private key. A verification procedure can use the security value to detect unauthorised changes to the assembly, but we do not

---

<sup>2</sup>The CLI specification allows resolution to take place when an application is installed, but we do not know of any implementation that takes such an eager approach.

consider verification further since it happens after binding. However, whether or not an assembly name is strong significantly affects binding, so this definition will be useful:

$$\text{StrongName}(\alpha) \equiv \text{Key}(\alpha) \neq \epsilon$$

An *assembly definition* consists of an assembly name, assembly dependencies and class definitions. Bytecode refers to assemblies by their display name, so the dependencies map display names to full assembly names. We assume that bytecode is encapsulated in class definitions of some type *ClassDef*. An assembly definition knows the location of the file that it was loaded from; this is used in type-casting and reflection operations.

$$\text{AssemblyDefinition } \delta : AD =$$

$$\text{Name} : AN, \text{Refs} : id \longrightarrow AN, \text{Code} : id \longrightarrow \text{ClassDef}, \text{Loc} : id$$

The CLI specification defines an assembly as comprised of modules (that contain bytecode) and other resource files. An assembly's module and resource files may be placed in a single physical file or left as independent files. However, modules and the physical layout of an assembly play no role in binding<sup>3</sup>, so we ignore them in our model. This keeps the definition of the *Code* element simple.

## 2.2 Assembly environment

Most CLI implementations (though not the CLI specification itself) support a standard location on disk where assemblies can be placed, typically if they have a strong name. At load-time, this location is typically checked before others, and thus provides the default environment from which assemblies come. In the CLR and Mono, the environment is provided by the Global Assembly Cache (GAC).

$$\text{Environment } \Delta : Env = AN \longrightarrow AD$$

We introduce the *extended environment* to represent both the filesystem of the machine executing the code, and a URL-addressable space of machines that have assemblies available for download. Given a list of paths, the extended environment tries each in turn until an assembly definition is found; it returns  $\epsilon$  if the list is exhausted without finding an assembly.

$$\text{Extended Environment } EE : \text{ExtEnv} = id^* \longrightarrow AD$$

## 3 Assembly-Oriented Execution

### 3.1 State

We wish to show how assembly identity, resolution and loading affect execution. We distinguish the state of the executing program from the state of the runtime system itself. Program state  $P$  is a pair whose elements are an instruction stack and an operand stack.

$$\text{Program state } P : I^* \times V^*$$

A CLI instruction  $I$  is parameterised by a display name and a member descriptor,  $M$ . The display name must have a corresponding entry in the *Refs* element of the enclosing assembly. A member descriptor is simply a class and a field/method signature. Values come from a type  $V$  with which we are not concerned.

The runtime system's state is represented by three elements: an environment (defined in section 2), a heap and a stack.

The runtime system's heap stores assembly definitions loaded from the environment and extended environment. The CLR's heap is divided into two parts, called *contexts*.<sup>4</sup> Contexts stop

<sup>3</sup>Partition I, §9.6: "... rather than establishing relationships between individual modules and referenced assemblies, every reference is resolved through the current assembly. This allows each assembly to have absolute control over how references are resolved."

<sup>4</sup>In fact, there is a third heap context, but its role is not important in our current model.

a programmer circumventing the system’s binding policies. The Mono system does not support contexts at present.

An assembly loaded by the CLR itself is placed in the first context. This happens when a bytecode instruction is jitted and the instruction’s display name is resolved. Assemblies loaded directly from a filename are placed in the second context. This happens when a programmer uses the reflection API provided by the core assemblies in the CLR and Mono. With a heap consisting of a pair of mappings from assembly name to definition, we write  $H_x$  for  $H \downarrow_1$ .

$$\text{Heap } H : (AN \longrightarrow AD) \times (AN \longrightarrow AD)$$

In a CLI implementation, the heap of loaded assemblies is part of an appdomain, which is a logical unit of isolation in a process. As we do not model the ability of a program to dynamically create and destroy appdomains, there is exactly one appdomain per executing application. Therefore, we do not need to qualify our heap of assemblies with an appdomain.

We need to track the call stack of assemblies at each dynamic program point. This is because the references of the currently executing assembly is consulted when resolving a reference to another assembly.<sup>5</sup> In addition, the context of the currently executing assembly is important when resolving an assembly reference. The stack starts with the assembly that the operating system considers is the entrypoint for an application.

$$\text{Stack } S : (AN \times \{1, 2\})^*$$

## 3.2 Evaluation

Evaluation is performed by a small-step operational semantics that evolves the state of the runtime system  $(\delta, H, S)$  and the program state  $(P)$ .

$$\Delta, H, S, P \longrightarrow \Delta', H', S', P'$$

The rules are shown in fig. 1. The bytecode instruction on the program’s instruction stack can be evaluated if it depends on an assembly already loaded into the system heap. (Rules EXEC-INSTR, EXEC-INSTR-CALL, EXEC-INSTR-CAST) Details of the evaluation are not important, so we abstract it into this judgement which evolves the program state given an assembly definition needed by the instruction:

$$\delta, P \longrightarrow P'$$

We are forced to differentiate the call instruction from other instructions because we need to add the called assembly’s name to the system stack, and modify the program’s instruction stream with the body of the called method. We assume a *lookup* function that can find a member  $M$  in an assembly.

A binding step can take place to resolve and load an assembly that an instruction is dependent on. (Rule EXEC-BIND) It uses the binding rules that evolve an environment and heap with an assembly definition for assembly name  $\alpha$ , returning the name of the actual assembly loaded:

$$\Delta, H, \alpha \longrightarrow \Delta', H', \alpha'$$

The execution is stuck if binding fails to find an assembly definition, *i.e.*  $\alpha'$  is  $\epsilon$ .

### 3.2.1 Heap contexts in evaluation

To evaluate a bytecode instruction, a definition must be available for the assembly it refers to. As per the CLI specification, we take the display name  $N$  mentioned in an instruction and look it up in the references of the currently executing assembly  $T$ , obtaining a full assembly name  $\alpha$ . In the CLR, which heap context to look up this assembly name  $\alpha$  in depends on which context

---

<sup>5</sup>CLI Specification Partition 1 §9.6

the currently executing assembly is loaded in. An assembly loaded in the first context can only “see” assemblies also loaded in the first context; an assembly loaded in the second context can see assemblies in both contexts, preferring the second. This policy is justified by the first context being where assemblies are “officially” loaded and the second context being where expert programmers place their own assemblies. (Mono only has one context, so the issue does not arise.)

$$\text{context}^{CLR}(\alpha, H, x) = \begin{cases} x & \text{if } \alpha \in \text{dom}(H_x) \\ 1 & \text{if } x = 2 \wedge \alpha \notin \text{dom}(H_2), \in \text{dom}(H_1) \end{cases}$$

$$\text{context}^{Mono}(\alpha, H, x) = x$$

### 3.2.2 Casting

Casting is complicated because assemblies play the same role as classloaders in Java, *i.e.* scoping a class such that a type is an (assembly name, class name) pair. Ensuring that the same classes from different assemblies are not confused is an important defence against attacks. Therefore, in the CLR, the source and target classes must be defined in the same *assembly file on disk*.

In addition, the heap context in which an assembly is loaded provides another level of qualification for a class, *i.e.* a type in the CLR is a (context id, assembly name, class name) triple. The same assembly definition can be loaded into multiple contexts, but casting an object across contexts would give rise to the same problems as casting it across classloaders. Therefore, the assembly definitions containing the source and target classes must be in the same context.

The EXEC-INSTR-CAST rule first obtains the full assembly name  $\alpha$  referred to by the `castclass` instruction. We assume that the object to be cast is accessible via the top value  $v$  on the program state’s value stack, and that the auxiliary function `type` returns an (assembly name, class name) pair representing the object’s type. The assemblies named by  $\alpha$  and  $\alpha'$  must be loaded, potentially in different contexts. We check that the two loaded assemblies were loaded from identical paths, as required by the CLR. If so, then the success of the cast is for the program to determine; we assume a notional `cast` operator that checks subclassing using the class definitions provided from an assembly definition:

$$\text{Code}(H_y(\alpha)), P[\text{cast } C' \text{ to } C] \longrightarrow P'$$

## 4 Assembly Binding

The binding rules in fig. 2 take a logical assembly name and return an assembly definition plus a name. If the assembly is not already loaded in the heap, then they use a name resolver  $\eta$ , a location resolver  $\odot$ , an assembly installer  $\oplus$ , and a name matcher  $\sim$ .

A name resolver performs a logical-to-logical mapping, applying versioning policy to an assembly name in order to obtain a more refined assembly name. A location resolver performs a logical-to-physical mapping, taking an assembly name and applying a “probing” policy that describes where to search for an assembly definition. If the location resolver fails to provide a location where a suitable assembly can be found, then an on-demand (*i.e.* “just-in-time”) assembly installation operation is tried, via  $\oplus$ .

If the extended environment is able to find an assembly, or an assembly is installed on-demand, then the binding rules return the heap augmented with the assembly definition, plus the name of the assembly that was actually loaded. CLI implementations require that the loaded name matches the name of the desired assembly (*i.e.* produced by the name resolver), according to  $\sim$ .

$$\text{Name Resolver } \eta : AN \longrightarrow AN$$

$$\text{Location Resolver } \odot : AN \times E \longrightarrow id^*$$

$$\text{Installer } \oplus : AN \times E \longrightarrow E$$

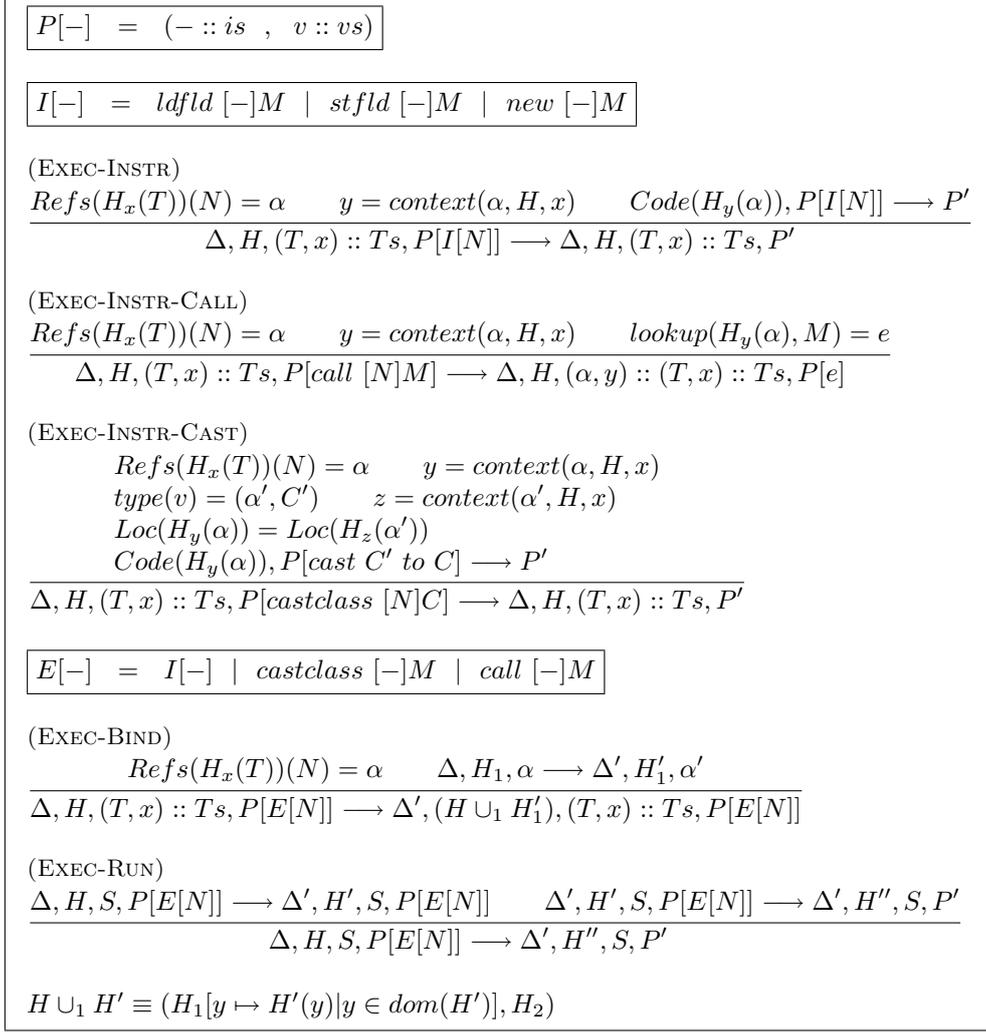


Fig. 1: Execution and Loading

*Name Matcher*  $\sim: AN \times AN \longrightarrow bool$

We introduce an *application context* that stores facts about the runtime environment for use by the name and location resolvers.

*Application Context*  $\Gamma$  :

(*RuntimeVersion* :  $int \times int \times int \times int$ , *Mapping* :  $AN \longrightarrow (AN \times id)$ , *AppBase* :  $id$ )

We define a Binding Framework  $BF = (\Gamma, \Delta, \eta, \odot, \oplus, \sim)$ . A binding framework is instantiated for a specific machine, CLI implementation and application. The local disk of the machine supplies the environment  $\Delta$ , which is a single directory for the CLR and one or more directories for Mono. The CLI implementation supplies  $\Gamma_{RuntimeVersion}$ ,  $\eta$ ,  $\odot$  and  $\oplus$ . The executed application supplies  $\Gamma_{AppPath}$  and  $\Gamma_{Mapping}$ . This models “side-by-side execution”, where several CLI implementations can be installed on the same machine, each with its own core assemblies, and the operating system chooses which implementation is suitable for executing a given application.

## 4.1 Name Resolution

A name resolver  $\eta$  maps a logical assembly name to another logical assembly name, according to three policies: servicing, unification and retargeting. Fig. 3 shows name resolvers for the CLR

$\frac{\text{(BIND-ALREADY-LOADED)} \quad \alpha \in \text{dom}(H)}{\Delta, H, \alpha \longrightarrow \Delta, H, \alpha}$	$\frac{\text{(BIND-AVAILABLE)} \quad \alpha \notin \text{dom}(H) \quad \eta(\alpha) = \alpha' \quad EE(\Delta \odot \alpha') = \delta \quad \alpha' \sim \text{Name}(\delta)}{\Delta, H, \alpha \longrightarrow \Delta, H[\alpha \mapsto \delta], \text{Name}(\delta)}$
$\frac{\text{(BIND-INSTALL-ON-DEMAND)} \quad \alpha \notin \text{dom}(H) \quad \eta(\alpha) = \alpha' \quad EE(\Delta \odot \alpha') = \epsilon \quad \Delta \oplus \alpha' = \Delta' \quad \alpha' \sim \text{Name}(\Delta'(\alpha))}{\Delta, H, \alpha \longrightarrow \Delta', H[\alpha \mapsto \Delta'(\alpha)], \text{Name}(\Delta'(\alpha))}$	
$\frac{\text{(BIND-UNAVAILABLE)} \quad \alpha \notin \text{dom}(H) \quad \eta(\alpha) = \alpha' \quad EE(\Delta \odot \alpha') = \epsilon \quad \Delta \oplus \alpha' = \Delta}{\Delta, H, \alpha \longrightarrow \Delta, H, \epsilon}$	

Fig. 2: Binding

and Mono.

**Servicing policy** To allow assemblies to be *serviced* (*i.e.* upgraded for security and performance reasons without modifying calling applications), the CLR supports policies for redirecting references to strongly-named assemblies. A reference to a non-strongly-named assembly cannot be serviced. (Mono does not currently support redirection policies.)

First, each application can supply a policy file for redirecting one version of a given assembly to another. Second, “publisher policies” can redirect requests for assemblies in the GAC. Third, a machine-wide redirection policy is applied after the application and publisher policies. We represent the union of these policies as a mapping from assembly name to assembly name in  $\Gamma_{\text{Mapping}}$  (using the first element of the range).

**Unification policy** A CLI-compliant virtual machine, such as the CLR, is often developed by different individuals from those who program the core assemblies that accompany the VM.<sup>6</sup> It is often practical to test a VM only with the exact framework assemblies that will accompany it.

The CLR and Mono both impose a restriction that some core assemblies (the exact set differs) must be the same version as that of the runtime execution system itself.

**Retargeting policy** As well as the CLR, Microsoft produces a CLI implementation for mobile devices called the .NET Compact Framework. An application compiled *for the CLR* will not run on a mobile device equipped with just the .NET Compact Framework, even if the developer is careful to use only assemblies available in the Compact Framework. This is because the core assemblies that accompany the CLR have different strong names from the assemblies in the Compact Framework [Mot04].

However, an application compiled *for the .NET Compact Framework* will run on the CLR. This is possible because the generated assembly references the Compact Framework’s assemblies by their strong names, as usual, but each reference features a *retargetable* flag. The .NET Compact Framework’s runtime ignores this flag and resolves the core assemblies as usual. The CLR reacts to it by rewriting the retargetable assembly names to the relevant core assembly names; the version number is unified and the key token is set to a standard value that indicates a core assembly to Fusion. This is Microsoft-specific behaviour; the Mono runtime will halt on failing to resolve the strong names of the Compact Framework assemblies referenced by the application.

<sup>6</sup>In Java, the `java.lang.*` class hierarchy.

$$\begin{array}{l}
\eta^{CLR}(\alpha) = \begin{cases} \alpha \\ \text{if } \neg \text{StrongName}(\alpha) \\ \\ \Gamma_{Mapping}(\alpha) \downarrow_1 \\ \text{if } \text{StrongName}(\alpha) \wedge \neg \text{Core}^{CLR}(\alpha) \\ \\ \alpha[\text{Version} \mapsto \Gamma_{RuntimeVersion}] \\ \text{if } \text{StrongName}(\alpha) \wedge \text{Core}^{CLR}(\alpha) \wedge \neg \text{Retargetable}(\alpha) \\ \\ \alpha[\text{Version} \mapsto \Gamma_{RuntimeVersion}, \text{Security} \mapsto \text{'b77a5c561934e089'}] \\ \text{if } \text{StrongName}(\alpha) \wedge \text{Core}^{CLR}(\alpha) \wedge \text{Retargetable}(\alpha) \end{cases} \\
\eta^{Mono}(\alpha) = \begin{cases} \alpha \\ \text{if } \neg \text{Core}^{Mono}(\alpha) \\ \\ \alpha[\text{Version} \mapsto \Gamma_{RuntimeVersion}] \\ \text{if } \text{Core}^{Mono}(\alpha) \end{cases} \\
\text{Core}^{CLR}(\alpha) \equiv \\ \quad \text{DisplayName}(\alpha) \in \{\text{mscorlib}, \text{System.Windows.Forms}, \dots\} \\
\text{Core}^{Mono}(\alpha) \equiv \text{StrongName}(\alpha) \wedge \text{DisplayName}(\alpha) \in \{\text{mscorlib}\}
\end{array}$$

Fig. 3: Name Resolution

## 4.2 Location resolution

A location resolver  $\odot$  supplies a list of physical filenames for the extended environment to try to obtain an assembly from. Fig. 4 shows location resolvers for the CLR and Mono.

Given an assembly name, the CLR’s location resolver prefers to search the environment first if the assembly’s name is a strongname. The next possible location is a “codebase” from the application context, specifically the second element of the  $\Gamma_{Mapping}$  entry for the target assembly name. The codebase’s location is final in the sense that no alternative paths are tried if it is specified. If a codebase is not specified, then various locations in the filesystem are suggested, using the path of the currently executing application (which is not necessarily that of the currently executing assembly). The extended environment will “probe” each of these locations in turn.

When performing location resolution for an assembly name that is not a strongname, the environment is not used. If a codebase is available, it must come from the same location as the executing application. Otherwise, the filesystem is tried as before.

The location resolver for Mono is quite different. It tries the application’s local directory first before the environment. (It also searches a CLASSPATH-style directory list before the environment, but we do not show this.)

## 4.3 Name matching

The CLR and Mono require an exact match between desired and loaded assembly versions:

$$\begin{array}{l}
\alpha \sim^{CLR, Mono} \alpha' \equiv \\
\text{StrongName}(\alpha) \iff \text{StrongName}(\alpha') \wedge \\
\text{Version}(\alpha) = a.b.c.d \iff \text{Version}(\alpha') = a.b.c.d
\end{array}$$

$$\Delta \odot^{CLR} \alpha = \begin{cases} \Delta, L & \text{if } StrongName(\alpha) \wedge \Gamma_{Mapping}(\alpha) \downarrow_2 = L \\ \Delta, Locs(\alpha) & \text{if } StrongName(\alpha) \wedge \Gamma_{Mapping}(\alpha) \downarrow_2 = \epsilon \\ L & \text{if } \neg StrongName(\alpha) \wedge \Gamma_{Mapping}(\alpha) \downarrow_2 = L \\ & \wedge L = \Gamma_{AppPath} + "/" + x \text{ for some } x \\ Locs(\alpha) & \text{if } \neg StrongName(\alpha) \wedge \Gamma_{Mapping}(\alpha) \downarrow_2 = \epsilon \end{cases}$$

$$Locs(\alpha) = (\Gamma_{AppPath} + "/" + DisplayName(\alpha) + ".dll"), \\
(\Gamma_{AppPath} + "/" + DisplayName(\alpha) + "/" + DisplayName(\alpha) + ".dll"), \\
(\Gamma_{AppPath} + "/" + Culture(\alpha) + "/" + DisplayName(\alpha) + ".dll"), \\
(\Gamma_{AppPath} + "/" + Culture(\alpha) + "/" + DisplayName(\alpha) + "/" + \\
\quad DisplayName(\alpha) + ".dll")$$

$$\Delta \odot^{Mono} \alpha = (\Gamma_{AppPath} + "/" + DisplayName(\alpha) + ".dll"), \Delta$$

Fig. 4: Location Resolution

$$\Delta \oplus^{CLR} \alpha = \Delta' \text{ for some } \Delta' \supseteq \Delta \text{ where } \Delta'(\alpha) \neq \epsilon \implies \Delta(\alpha) = \epsilon$$

$$\Delta \oplus^{Mono} \alpha = \Delta$$

Fig. 5: Software Installation

#### 4.4 Install-on-demand

If both the environment and the extended environment fail to supply an assembly, the  $\oplus$  function tries to perform an “install-on-demand” operation. Unlike the extended environment, which is queried at a specific location (*e.g.* a URL), the installer is required to return an assembly given just its name.

In the CLR, we suppose that the end user is asked to supply an assembly, *e.g.* on a CD. Because the supplied assembly is totally free, we pass the old environment to  $\oplus$  to see that if it does grow, then a truly new assembly is available in the new environment. This approach allows us to accept that the installation can fail, leaving the environment unchanged and propagating (through binding rule BIND-UNAVAILABLE) a loading failure.

Mono does not support on-demand installation, so returns an unchanged environment.

#### 4.5 Dynamic loading through reflection

As stated in section 2, assemblies can be loaded using a reflection API. This is widely used by developers building applications that support plug-ins. Among the many reflection methods provided by the CLR’s core assemblies, we consider `Load` and `LoadFrom`. Mono’s core assemblies provide `Load` only. The full method signatures are shown in fig. 6.

`Load` takes a display name  $N$  from the program state’s value stack, and defers to the runtime’s binding rules to resolve and load it. The behaviour here is exactly that of the EXEC-LOAD rule in fig. 1. (A CLI implementation will typically also provided an overloaded `Load` that takes a full assembly name, but this merely allows lookup in *Refs* to be avoided.)

`LoadFrom` is more complex. It takes a location  $L$  from the program state's value stack, and loads the file at that location directly. It then performs what Microsoft call a *second bind* by taking the *display name* of the just-loaded assembly definition and trying to bind it. If binding gives the same assembly (by physical file location) as one already loaded, then the heap is updated with the result of the *second bind*. If the second bind cannot find an assembly, or it finds a different physical file to that directly loaded, then the directly-loaded assembly is placed in the heap's second context.

$$\begin{array}{l}
LOAD \equiv call[mscorlib]System.Reflection.Assembly :: Load \\
LOADFROM \equiv call[mscorlib]System.Reflection.Assembly :: LoadFrom \\
\\
\boxed{P[-, -] = (- :: is, - :: vs)} \\
\\
(\text{EXEC-INSTR-CALLLOAD}) \\
\frac{Refs(H_x(T))(N) = \alpha \quad \Delta, H_1, \alpha \longrightarrow \Delta', H'_1, \alpha'}{\Delta, H, (T, x) :: Ts, P[LOAD, N] \Longrightarrow \Delta', (H \cup_1 H'_1), (T, x) :: Ts, P[\epsilon, \alpha']} \\
\\
(\text{EXEC-INSTR-CALLLOADFROM}) \\
\frac{EE(L) = \delta \quad \Delta, H_1, DisplayName(Name(\delta)) \Longrightarrow \Delta', H'_1, \alpha \\
\alpha \neq \epsilon \wedge Loc(\delta) = Loc(H'_1(\alpha))}{\Delta, H, S, P[LOADFROM, L] \Longrightarrow \Delta', (H \cup_1 H'_1), S, P[\epsilon, \alpha]} \\
\\
(\text{EXEC-INSTR-CALLLOADFROM2}) \\
\frac{EE(L) = \delta \quad \Delta, H_1, DisplayName(Name(\delta)) \Longrightarrow \Delta', H'_1, \alpha \\
\alpha = \epsilon \vee Loc(\delta) \neq Loc(H'_1(\alpha))}{\Delta, H, S, P[LOADFROM, L] \Longrightarrow \Delta, (H \cup_2 [Name(\delta) \mapsto \delta]), S, P[\epsilon, Name(\delta)]} \\
\\
H \cup_1 H' \text{ as before} \\
H \cup_2 [\alpha \mapsto \delta] \equiv (H_1, H_2[\alpha \mapsto \delta | \alpha \notin dom(H_2)])
\end{array}$$

Fig. 6: Dynamic loading through reflection

## 5 Related & Further Work

Classloading in Java has received significant attention [Dea97, JLT98, LB98, QGC00], and [FZA04] presents it in an abstract setting. However, relatively little work focuses on the CLI platform. [DLE03] unifies dynamic linking in Java and the CLI, but abstracts the assembly binding process to a very high level. [EJS02] and [EJS03] offer a formal model of a well-formed GAC, where assembly addition and removal do not break existing dependencies. Our work is clearly complementary to this, as we show how the GAC is used in the wider assembly binding process. Our  $\oplus$  operator would ideally maintain a stronger safety property concerning evolution of the GAC [EJS03].

We have described and formalised how assemblies are resolved and loaded by common CLI implementations. Most programmers assume that an assembly's strong name is its sole identity once loaded, but we show how the CLR, during execution, considers an assembly's identity to have more elements. Namely, it considers where an assembly was loaded from (*i.e.* a disk or URL-based location) and where it was loaded to (*i.e.* its heap context). These elements are necessary because the CLR exposes reflective assembly loading operations that can load arbitrary assemblies. While merely loading such assemblies is harmless, it is essential to avoid using their classes if the assembly's identity masquerades as one of the core assemblies. We plan to state formally that binding is "safe" in the current CLR in that it never leads to a heap where a non-core assembly

is mistaken for a core assembly. The Mono system avoids the problem at present by not offering reflective loading capabilities.

A weakness of the current model is that name resolution produces a very precise answer, *i.e.* a single assembly name. This does not accurately model the .NET Compact Framework or, indeed, more flexible future schemes for choosing an assembly to load [BD04, BMED05]. The .NET Compact Framework does not support servicing policies that redirect an assembly's desired version, so applications cannot be directed to use later, better code. However, the Compact Framework's binding rules permit the loader to provide version  $a.b.c.x$  of an assembly when a reference is made to version  $a.b.c.d$ , *i.e.* the last element of the version number can “float”. The binding rules also permit *any* version of an assembly to be loaded when the reference mentions version 0.0.0.0.

In our model, this equates to the name resolver producing  $a.b.c.*$  for the desired version to locate. We could modify name resolution to produce a constraint on permitted names, rather than a specific name. Location resolution would then need to iterate through the files found in the extended environment to choose the “best” one matching the constraint. The name matcher would have the following definition:

$$\alpha \sim_{CompactFramework} \alpha' \equiv \\ StrongName(\alpha) \iff StrongName(\alpha') \wedge \\ ((Version(\alpha) = a.b.c._ \iff Version(\alpha') = a.b.c._) \vee Version(\alpha) = 0.0.0.0)$$

The CLR v2.0 will be released in late 2005 and makes some small changes to unification policy [Shi05], so we will need a new name resolver. More interesting are Microsoft's plans for binding in Longhorn [GR04], where assemblies are typed and servicing policy is affected by the types of referencing and referenced assemblies. A feature called “interim roll-back” is also planned, where assemblies installed in the environment are temporarily hidden due to flaws being found in them. Our model can handle the new servicing policy (at name resolution) and rollback policy (at location resolution). More challenging is to state whether syntactic or semantic compatibility is assured by these new features.

## References

- [BD04] Alex Buckley and Sophia Drossopoulou. Flexible Dynamic Linking. In *ECOOP Workshop on Formal Techniques for Java Programs (FTfJP 2004)*, Oslo, Norway, June 2004.
- [BMED05] Alex Buckley, Michelle Murray, Susan Eisenbach, and Sophia Drossopoulou. Flexible Bytecode for Linking in .NET. In *First Workshop on Bytecode Semantics, Verification, Analysis and Transformation (BYTECODE 2005)*, ENTCS, Edinburgh, Scotland, March 2005. Elsevier BV.
- [Coo05] Suzanne Cook. .NET CLR Loader Notes. <http://blogs.msdn.com/suzcook>, 2005.
- [Dea97] Drew Dean. The Security of Static Typing with Dynamic Linking. In *Proceedings of the Fourth ACM Conference on Computer and Communications Security*, Zurich, Switzerland, April 1997.
- [dI05] Miguel de Icaza. Mono. <http://www.mono-project.com/>, 2005.
- [DLE03] Sophia Drossopoulou, Giovanni Lagorio, and Susan Eisenbach. Flexible Models for Dynamic Linking. In Pierpaolo Degano, editor, *Proceedings of the 12th European Symposium on Programming (ESOP 2003)*, volume 2618 of *LNCS*, pages 38–53. Springer-Verlag, April 2003.
- [ECM02] ECMA. *Standard ECMA-335: Common Language Infrastructure*. ECMA International, December 2002. <http://www.ecma-international.org/publications/standards/Ecma-335.htm>.
- [EJS02] S. Eisenbach, V. Jurisic, and C. Sadler. Feeling the way through DLL Hell. In *Proceedings of the First Workshop on Unanticipated Software Evolution (USE 2002)*, Malaga, Spain, June 2002.
- [EJS03] S. Eisenbach, V. Jurisic, and C. Sadler. Managing the Evolution of .NET Programs. In *6th IFIP International Conference on Formal Methods for Open Object-based Distributed Systems (FMOODS 2003)*, volume 2884 of *LNCS*, pages 185–198, Paris, France, November 2003. Springer-Verlag.
- [Far05] Shawn Farkas. .NET Security Blog. <http://blogs.msdn.com/shawnfa>, 2005.
- [FZA04] Sonia Fagorzi, Elena Zucca, and Davide Ancona. Modeling Multiple Class Loaders by a Calculus for Dynamic Linking. In *Proceedings of the ACM Symposium on Applied Computing (SAC-2004)*, Nicosia, Cyprus, March 2004.

- [GR04] Cathi Gero and Jeffrey Richter. The Future of Assembly Versioning. <http://www.theserverside.net/articles/showarticle.tss?id=AssemblyVersioning>, 2004.
- [JLT98] T. Jensen, D. Le Metayer, and T. Thorn. Security and Dynamic Class Loading in Java: A Formalisation. In *Proceedings of the IEEE International Conference on Computer Languages*, pages 4–15, Chicago, IL, USA, 1998.
- [LB98] Sheng Liang and Gilad Bracha. Dynamic Class Loading in the Java Virtual Machine. In *Proceedings of the 13th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA'98)*, Vancouver, BC, Canada, October 1998.
- [MG00] Eric Meijer and John Gough. *Technical Overview of the Common Language Runtime*. Microsoft, 2000.
- [Mot04] Daniel Moth. <http://www.danielmoth.com/Blog>, 2004.
- [Pra05] Steven Pratschner. .NET CF WebLog. <http://blogs.msdn.com/stevenpr>, 2005.
- [QGC00] Zhenyu Qian, Allen Goldberg, and Alessandro Coglio. A Formal Specification of Java Class Loading. In *Proceedings of the 15th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA 2000)*, pages 325–336, Minneapolis, MN, USA, 2000.
- [Shi05] Alan Shi. The Fusion Weblog. <http://blogs.msdn.com/alanshi>, 2005.
- [Zha05] Junfeng Zhang. .NET Framework Notes. <http://blogs.msdn.com/junfeng>, 2005.

## A Examples

### A.1 Context-locality of dependencies

We produce two strongly named assemblies, `Server` and `Server2`. Code in `Server` calls code in `Server2`. We wish to show that if both assemblies are loaded in the second context, then a dependency from one to the other is automatically resolved.

`Server2` is built first (producing `Server2.dll`), then `Server` is built (referencing `Server2.dll` and producing `Server.dll`). `Server.dll` and `Server2.dll` exist in separate directories, both unknown to Fusion.

Listing 1: `Server` assembly

```
1 using System;
2 using System.Reflection;
3
4 [assembly: AssemblyVersion("1.0.0.0")]
5 [assembly: AssemblyKeyFile("Server.key")]
6
7 public class Server {
8     public void UseServer2() {
9         Console.WriteLine("Calling other server");
10        Server2 other = new Server2();
11        other.print();
12    }
13 }
```

Listing 2: `Server2` assembly

```
1 using System;
2 using System.Reflection;
3
4 [assembly: AssemblyVersion("1.0.0.0")]
5 [assembly: AssemblyKeyFile("Server2.key")]
6
7 public class Server2 {
8     public void print() {
9         Console.WriteLine("I am Server2");
10    }
11 }
```

The Client class in listing 3 is placed in an assembly that has no static dependencies on the `Server` or `Server2` assemblies. The `LoadFrom` statements place both these assemblies in the second context. The different hash codes in the output indicate that the different `.dll` files are distinct Assembly objects in memory.

Listing 3: Dependencies are context-local

```
1 // Dependencies are context-local
2
3 using System;
4 using System.Reflection;
5
6 public class Client {
7     public static void Main(String[] args) {
8         Assembly a = Assembly.LoadFrom(@"c:\scenario1\hidden1\Server.dll");
9         Console.WriteLine(a+" at "+a.Location);
10        Console.WriteLine(a.GetHashCode());
11    }
```

```

12     Assembly b = Assembly.LoadFrom(@"c:\scenario1\hidden2\Server2.dll");
13 //     Assembly b = Assembly.Load("Server2, Version=1.0.0.0, Culture=neutral
14     , PublicKeyToken=4223ED1A651AD56B");
15     Console.WriteLine(b+" at "+b.Location);
16     Console.WriteLine(b.GetHashCode());
17
18     Object o = a.CreateInstance("Server");
19     Type t = o.GetType();
20     MethodInfo m = t.GetMethod("UseServer2");
21     m.Invoke(o, null);
22 }

```

By indirect invocation at line 20 of the client, code in the Server code is invoked. (We need to use indirect invocation to avoid a static reference to the Server assembly from the Client. Since the Client executable is in the first context, having been loaded by the operating system, it would not be able to use the Server assembly in the second context.)

Now, code in the Server assembly (class Server, method UseServer2) calls code in the Server2 assembly. The static dependency that Server has on Server2 is successfully resolved *within the second context*. No attempt to load a Server2 assembly takes place.

Listing 4: Output of listing 3

```

1 C:\scenario1>client
2 Server , Version=1.0.0.0 , Culture=neutral , PublicKeyToken=fc71b27c35bb8a9d at
   c:\scenario1\hidden1\server.dll
3 2
4 Server2 , Version=1.0.0.0 , Culture=neutral , PublicKeyToken=4223ed1a651ad56b
   at c:\scenario1\hidden2\server2.dll
5 3
6 Calling other server
7 I am Server2

```

Note that if we comment out lines 12 and uncomment line 13 in the Client class, then we see *cross-context* dependency resolution. Assuming that Server2.dll is able to be found (*e.g.* it is in the same directory as the Client executable), it will be loaded into the first context. The code in the Server assembly, loaded in the second context, *is* able to call code in the Server2 assembly, loaded in the first context.

## A.2 Unification

This example shows that trying to bind to an arbitrary core assembly results in the bind succeeding, but to the “approved” core assembly.

This program tries to load some unavailable versions of the `mscorlib` core assembly.

Listing 5: Unification

```
1 // References to mscorlib are automatically unified
2
3 using System;
4 using System.Reflection;
5
6 public class Client {
7     public static void Main(String[] args) {
8         Assembly a = Assembly.Load("Mscorlib, Version=2.0.0.0, Culture=neutral,
9             PublicKeyToken=B77A5C561934E089");
10        Console.WriteLine(a+" at "+a.Location);
11        Assembly b = Assembly.Load("Mscorlib, Version=0.0.0.1, Culture=neutral,
12            PublicKeyToken=B77A5C561934E089");
13        Console.WriteLine(b+" at "+b.Location);
14    }
15 }
```

Running it on the CLR v1.1.4322.0, we see unification happening, so only the correctly versioned assembly is loaded.

```
1 C:\scenario2>client
2 mscorlib, Version=1.0.5000.0, Culture=neutral, PublicKeyToken=
   b77a5c561934e089 at c:\windows\microsoft.net\framework\v1.1.4322\
   mscorlib.dll
3 mscorlib, Version=1.0.5000.0, Culture=neutral, PublicKeyToken=
   b77a5c561934e089 at c:\windows\microsoft.net\framework\v1.1.4322\
   mscorlib.dll
```

### A.3 Casting across assemblies

This example shows that casting an object from a class in one assembly, to the *same* class in another assembly, fails when the assemblies have different versions.

We create two assemblies, both called `Server.dll`, that differ *only* in their version number. We install both to the GAC.

Listing 6: Server assembly v1.0.0.0

```
1 using System;
2 using System.Reflection;
3
4 [assembly: AssemblyVersion("1.0.0.0")]
5 [assembly: AssemblyKeyFile("Server.key")]
6
7 public class Server {}
```

Listing 7: Server assembly v1.0.0.1

```
1 using System;
2 using System.Reflection;
3
4 [assembly: AssemblyVersion("1.0.0.1")]
5 [assembly: AssemblyKeyFile("Server.key")]
6
7 public class Server {}
```

The client loads both assemblies (into the first context), then tries to cast an instance of the `Server` class from the v1.0.0.1 assembly, to the `Server` class from the v1.0.0.0 assembly. The Client executable statically references v1.0.0.0 of the `Server` assembly, due to the instantiation of a `Server` object at line 8 in listing 8.

Listing 8: No casting between assembly versions

```
1 // A class isn't compatible with itself if it comes from different version
   of the same assembly
2
3 using System;
4 using System.Reflection;
5
6 public class Client {
7     public static void Main(String[] args) {
8         Server s = new Server();
9         Console.WriteLine("Statically-referenced Server class comes from
   assembly: "+s.GetType().Assembly.GetName()+" at "+s.GetType().
   Assembly.Location);
10
11         Assembly a = Assembly.Load("Server, Version=1.0.0.1, Culture=neutral,
   PublicKeyToken=17aa5269acb19216");
12         Object o = a.CreateInstance("Server");
13         Console.WriteLine("Dynamically-referenced Server class comes from "+o.
   GetType().Assembly+" at "+o.GetType().Assembly.Location);
14
15         // Try casting from the Server class in assembly 1.0.0.1 to the Server
   class in assembly 1.0.0.0
16         try {
17             Server s2 = (Server)o;
18         } catch (InvalidCastException ex) {
19             Console.WriteLine(ex);
```

```
20     }
21   }
22 }
```

The cast fails and a `InvalidCastException` is thrown. Even though the `Server` class itself is identical in both the `Server` and `Server2` assemblies, the fact that the assemblies have different versions is important.

```
1 C:\scenario3>client
2 Statically-referenced Server class comes from assembly: Server, Version
   =1.0.0.0, Culture=neutral, PublicKeyToken=17aa5269acb19216 at c:\windows
   \assembly\gac\server\1.0.0.0__17aa5269acb19216\server.dll
3 Dynamically-referenced Server class comes from Server, Version=1.0.0.1,
   Culture=neutral, PublicKeyToken=17aa5269acb19216 at c:\windows\assembly\
   gac\server\1.0.0.1__17aa5269acb19216\server.dll
4 System.InvalidCastException: Specified cast is not valid.
5   at Client.Main(String[] args)
```

## A.4 Casting across contexts

This example shows that casting an object from a class in one assembly, to the *same* class in another assembly, fails when the assemblies are loaded in different contexts.

We use the v1.0.0.0 Server assembly from scenario A.3, and build Server.dll. We install it into the GAC, and also copy it to a directory unavailable to Fusion.

Listing 9: Server assembly v1.0.0.0

```
1 using System;
2 using System.Reflection;
3
4 [assembly: AssemblyVersion("1.0.0.0")]
5 [assembly: AssemblyKeyFile("Server.key")]
6
7 public class Server {}
```

The client references the Server class, so when we build the client executable, we must reference the “hidden” Server.dll that provides the class. This gives the client executable a static reference to an assembly called Server, of version v1.0.0.0. Such a reference causes the server to load the Server.dll assembly in the GAC. This assembly goes into the first context. The client also loads the “hidden” Server.dll via LoadFrom, so the assembly goes into the second context.

The client then tries to cast an instance of the Server class from the Server.dll loaded from the filesystem, to the class from the Server.dll loaded from the GAC.

Listing 10: No casting across contexts

```
1 // A class isn't compatible with itself if it comes from assemblies in
   // different contexts (hence their paths are different)
2
3 using System;
4 using System.Reflection;
5
6 public class Client {
7     public static void Main(String[] args) {
8         Server s = new Server();
9         Console.WriteLine("Statically-referenced Server class comes from
   // assembly: "+s.GetType().Assembly+" at "+s.GetType().Assembly.
   // Location);
10
11         Assembly a = Assembly.LoadFrom(@"C:\scenario4\hidden\Server.dll");
12         Object o = a.CreateInstance("Server");
13         Console.WriteLine("Dynamically-referenced Server class comes from
   // assembly: "+o.GetType().Assembly+" at "+o.GetType().Assembly.
   // Location);
14
15         // Try casting from the Server class in the LoadFrom context to the
   // Server class in the Load context
16         try {
17             Server s2 = (Server)o;
18         } catch (InvalidCastException ex) {
19             Console.WriteLine(ex);
20         }
21     }
22 }
```

Since the assemblies that declare Server classes are in different contexts, the cast fails, even though the Server classes are identical. (Indeed, the assemblies themselves are identical; only their location, and hence context, is different.)

```
1 C:\scenario4>client
2 Statically-referenced Server class comes from assembly: Server, Version
   =1.0.0.0, Culture=neutral, PublicKeyToken=17aa5269acb19216 at c:\windows
   \assembly\gac\server\1.0.0.0__17aa5269acb19216\server.dll
3 Dynamically-referenced Server class comes from assembly: Server, Version
   =1.0.0.0, Culture=neutral, PublicKeyToken=17aa5269acb19216 at c:\
   scenario4\hidden\server.dll
4 System.InvalidCastException: Specified cast is not valid.
5   at Client.Main(String[] args)
```

## A.5 LoadFrom reuses assemblies

This example shows that `LoadFrom` has the surprising behaviour that it does not always load the assembly specified.

We use the v1.0.0.0 Server assembly from scenario A.3. We copy it to two different locations, both unavailable to Fusion. The client then uses `LoadFrom` to attempt to load both files.

Listing 11: LoadFrom reuses assemblies

```
1 // In the LoadFrom context, loading the same assembly from different
  locations reuses the first one loaded
2
3 using System;
4 using System.Reflection;
5
6 public class Client {
7     public static void Main(String[] args) {
8         Assembly a = Assembly.LoadFrom(@"C:\scenario5\hidden1\Server.dll");
9         Console.WriteLine(a+" at "+a.Location);
10        Console.WriteLine(a.GetHashCode());
11
12        Assembly b = Assembly.LoadFrom(@"C:\scenario5\hidden2\Server.dll");
13        Console.WriteLine(b+" at "+b.Location);
14        Console.WriteLine(b.GetHashCode());
15
16        AppDomain currentDomain = AppDomain.CurrentDomain;
17        Assembly[] assems = currentDomain.GetAssemblies();
18        foreach (Assembly assem in assems) Console.WriteLine("Loaded: "+assem
19            +" "+assem.Location);
20    }
}
```

However, since the assembly's strong name is identical in both locations, the second `LoadFrom` merely receives the same `Assembly` object as the first (evidenced by the same hash code, 2).

For good measure, the client prints out all assemblies loaded in the appdomain, regardless of context. We see that only one Server assembly is loaded, namely the first.

```
1 C:\scenario5>client
2 Server, Version=1.0.0.0, Culture=neutral, PublicKeyToken=17aa5269acb19216 at
  c:\scenario5\hidden1\server.dll
3 2
4 Server, Version=1.0.0.0, Culture=neutral, PublicKeyToken=17aa5269acb19216 at
  c:\scenario5\hidden1\server.dll
5 2
6 Loaded: mscorlib, Version=1.0.5000.0, Culture=neutral, PublicKeyToken=
  b77a5c561934e089 c:\windows\microsoft.net\framework\v1.1.4322\mscorlib.
  dll
7 Loaded: Client, Version=0.0.0.0, Culture=neutral, PublicKeyToken=null C:\
  scenario5\Client.exe
8 Loaded: Server, Version=1.0.0.0, Culture=neutral, PublicKeyToken=17
  aa5269acb19216 c:\scenario5\hidden1\server.dll
```