

Object-oriented Programming in F#

Björn Lisper
School of Innovation, Design, and Engineering
Mälardalen University

`bjorn.lisper@mdh.se`
`http://www.idt.mdh.se/~blr/`

F# as an Object-oriented Programming Language

F# has an object-oriented part

Important to know something about this – accessing .NET libraries and services is done in an object-oriented fashion

F# has these kinds of *object types*:

- Concrete object types
- Object interface types

Concrete Object Types

F# objects can be both *mutable* and *immutable*

Methods are called *members*

The simplest case: extending a conventional F# type with *member declarations*

These provide a kind of interface to the data type

(They're really just functions that take an argument of that type)

An example

Extending a record data type with members, turning the records into objects:

```
type vector2D =  
    { x : float; y : float }  
    member v.Length = sqrt (v.x**2.0 + v.y**2.0)  
    member v.Scale(k) = { x = k*v.x; y = k*v.y }  
    member v.X_shift(x_new) = { v with x = x + x_new }  
    static member Zero = { x = 0.0; y = 0.0 }  
    static member X_vector(x_in) = { x = x_in; y = 0.0 }
```

`v` is like “`this`”, or “`self`” in other languages. In F# you can choose whatever identifier you want

Members without arguments are *properties*, members with arguments are *methods*

Members can be *static*, meaning they operate on types rather than values. We'll see an example soon

The above is an immutable object type: no record fields can be updated

This is the object type resulting from the declaration:

```
type vector2D =  
    { x : float; y : float }  
    member Length : float  
    member Scale : k:float -> vector2D  
    member X_shift : x_new:float -> vector2D  
    static member Zero : vector2D  
    static member X_vector : x_in:float -> vector2D
```

The member types become part of the object type

Creating an object and using its methods and properties:

```
let vec = { x = 3.0; y = 4.0 } // create an object "vec"
vec.Length                     // length of vec = 5.0
let vec2 = vec.Scale(2.0)      // create a new object "vec2"
vec2.Length                    // length of vec = 10.0
vec.Length                     // length of vec is still 5.0
vector2D.Zero                  // a new object with x, y = 0.0
```

Note the immutability: `vec.Scale(2.0)` creates a new object, the old one is not affected

Static members are applied to the type, not to values of that type

Members vs. Functions

Consider this alternative, F# with ordinary functions:

```
type vector2D =  
    { x : float; y : float }  
let Length v      = sqrt (v.x**2.0 + v.y**2.0)  
let Scale v k      = { x = k*v.x; y = k*v.y }  
let Zero          = { x = 0.0; y = 0.0 }  
let X_vector x_in = { x = x_in; y = 0.0 }
```

Perfectly possible, does the same thing. But we lose the bundling of members and record type into an object type

A Variation

Any F# type can be enriched with members into an object type:

```
type Tree<'a> = Leaf of 'a | Branch of Tree<'a> * Tree<'a>
    member t.Fringe =
        match t with
        | Leaf x -> [x]
        | Branch (t1,t2) -> t1.Fringe @ t2.Fringe
```

It doesn't have to be a record type

Members can be recursive

Constructed Classes

Goes beyond the simple object types where ordinary F# types are extended with members

Adds a possibility to define entities local to objects

These entities can be precomputed

An Example of a Constructed Class

vector2D using a constructed class:

```
type vector2D(x : float; y : float) =  
    let len = sqrt (x**2.0 + y**2.0)  
    member v.Length = len  
    member v.Scale(k) = vector2D(k*x, k*y)  
    member v.X_shift(x_new) = vector2D(x = x + x_new, y = y)  
    static member Zero = vector2D(x = 0.0, y = 0.0)  
    static member X_vector(x_in) = vector2D(x = x_in, y = 0.0)
```

`vector2D` is a *constructor* (in the OO sense): a function that creates a new object

```
let v = vector2D(3.0, 4.0)
```

`len` will be computed then `vector2D` creates the new object

Arguments to members can be given both by position (`Scale`), or by name (e.g., `X_shift`)

The resulting type:

```
type vector2D =  
    new : x:float * y:float -> Vector2D  
    member Length : float  
    member Scale k:float -> Vector2D  
    member X_shift x_new:float -> Vector2D  
    static member Zero : vector2D  
    static member X_vector : x_in:float -> Vector2D
```

Note “new”, tells the type of the `vector2D` constructor

Named and Optional Arguments

With named arguments, it is convenient to make arguments optional and have a default value for them

Named arguments can be used with all method calls

Optional arguments are preceded by “?”

An optional argument with type `'a` will have type `'a option` within the object type declaration

If the argument `v` is given, then it will have value `Some v` inside

If the argument is not given, it will have value `None`

It is the responsibility of the programmer to write code that uses this distinction to provide a default value

Optional Arguments, Example

We turn `x` and `y` into optional arguments with default `0.0`:

```
type vector2D(?x : float; ?y : float) =  
    let x = match x with  
        | None    -> 0.0  
        | Some v -> v  
    let y = match y with  
        | None    -> 0.0  
        | Some v -> v  
    .....
```

Note the new local definitions of `x` and `y` – not the same as the arguments `x` and `y`!

Optional Arguments, Continued

A builtin function to use with optional arguments:

```
defaultArg : 'a option -> 'a -> 'a
```

Its definition;

```
DefaultArg arg default =  
    match arg with  
    | None    -> default  
    | Some a  -> a
```

An example of its use:

```
defaultArg x 0.0
```

Mutable Object Types

One idea with object-orientation is to encapsulate side-effects into objects

This reduces the risks with the side-effects

Side-effects means we should have mutable data inside objects

F# supports this


Object-type internal variables can be declared `mutable`

Members are defined with `get` and `set` methods:

- The `get` method returns the current value for the member
- The `set` method sets a new value for the member by setting new values for the object-internal mutable variables

An Example

The 2D-vector again, but now with two different views:



An object representing a 2D-vector will have x and y as mutable state

However, we will also provide methods for *Length* and *Angle*

Getting *Length* and *Angle* will compute them from x and y

Setting *Length* or *Angle* will set x and y

Object Type Declaration

```
type mutVector2D(x : float; y : float) =  
    let mutable current_x = x  
    let mutable current_y = y  
    member v.x with get () = current_x and set x = current_x <- x  
    member v.y with get () = current_y and set y = current_y <- y  
    member v.Length  
        with get () = sqrt(current_x**2.0 + current_y**2.0)  
        and set len = let theta = v.Angle  
                        current_x <- len*cos theta  
                        current_y <- len*sin theta  
    member v.Angle  
        with get () = atan2 current_y current_x  
        and set theta = let len = v.Length  
                        current_x <- len*cos theta  
                        current_y <- len*sin theta
```

Resulting type:

```
type mutVector2D =  
    new : float * float -> mutVector2D  
    member x : float with get,set  
    member y : float with get,set  
    member Length : float with get,set  
    member Angle: float with get,set
```

How to Use

```
> let v = mutVector2D(3.0, 4.0);;
val v : mutVector2D
> (v.x, v.y);;
val it : float * float = (3.0, 4.0)
> (v.Length, v.Angle);;
val it : float * float = (5.0, 0.927295218)
> v.Length <- 10.0;;
val it : unit = ()
> (v.x, v.y);;
val it : float * float = (6.0, 8.0)
> (v.Length, v.Angle);;
val it : float * float = (10.0, 0.927295218)
> (v.x, v.y) <- (1.0, 1.0)
val it : unit = ()
> (v.Length, v.Angle);;
val it : float * float = (1.414213562, 0.7853981634)
```

Object Interface Types

“Abstract” object type declarations, specify only members and their types, not their implementations

Concrete implementations are specified by separate declarations

By having different concrete implementations implement members differently, we achieve something similar to virtual methods

Example

```
type Point = { X : float; Y : float }
type Rectangle = Rectangle of (float * float * float * float)

type IShape =
    abstract Contains : Point -> bool
    abstract Boundingbox : Rectangle

let circle(center:Point, radius:float) =
    { new IShape with
        member x.Contains(p:Point) =
            let dx = p.X - center.X
            let dy = p.Y - center.Y
            sqrt(dx**2.0 + dy**2.0) <= radius
        member x.Boundingbox =
            Rectangle (center.X - radius, center.Y - radius,
                       2.0*radius, 2.0*radius) }
```

```
let square(center:Point, side:float) =  
    { new IShape with  
        member x.Contains(p:Point) =  
            let dx = p.X - center.X  
            let dy = p.Y - center.Y  
            abs(dx) <= side/2.0 && abs(dy) <= side/2.0  
        member x.BoundingBox =  
            Rectangle (center.X - side,center.Y - side,  
                        2.0*side,2.0*side) }
```

Object Expressions

`circle` and `square` are functions whose bodies are *object expressions* (the “`{ new IShape with ... }`”)

Object expressions are used to specify implementations for interface types

An object expression must give an implementation for each member of the interface type

In our example, the functions `circle` and `square` provide implementations of the `IShape` interface

Inheritance

Object interface types can inherit from each other

Thus, hierarchies of such types can be built

The keyword “`inherit`” specifies inheritance

```
type Blahonga =  
    abstract xxx : ...
```

```
type FooBar =  
    inherit Blahonga  
    abstract yyy : ...
```

An implementation of `FooBar` must implement both `xxx` and `yyy`

Functional Programming Techniques and Object Expressions

We can define functions that return object expressions

In that way, object expressions can be abstracted

An example: a simple interface `TextOutputSink` defining two methods: for writing a character, and writing a string, and a function `SimpleOutputSink` returning an implementation

(See next page)

```
type TextOutputSink =  
    abstract WriteChar : char -> unit  
    abstract WriteString : string -> unit
```

```
let SimpleOutputSink(writechar) =  
    { new TextOutputSink with  
        member x.WriteChar(c) = writechar c  
        member x.WriteString(s) =  
            for c in s do writechar c }
```

SimpleOutputSink defines the simple pattern to write a string by writing it character by character

Using Objects in F#: A Simple GUI Example

Objects are needed in F# when interfacing with all the services of .NET

That's one reason why we spend time on the OO part of F# in the course

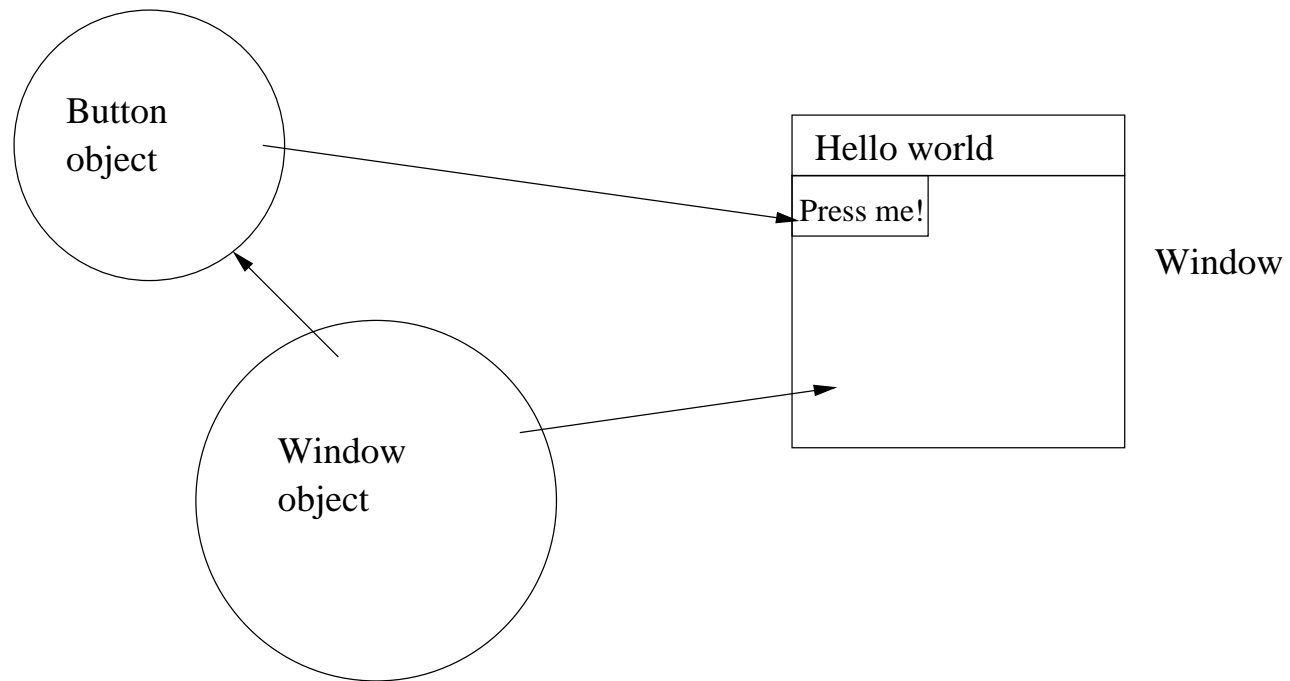
We'll bring up a simple example here: some simple GUI handling, with windows and buttons

Objects and GUI's

In F# (and .NET), each GUI component is represented by an object:

- A window
- A button
- Etc.

The object holds the representation: position, style, fill colour, text(s), etc.




Events

GUI's receive input from users: clicks, mouse movements, keys being pressed, ...

Represented by *events*


An event is basically a stream of data: coordinates for a mouse, data representing key clicks, etc.



GUI components (like a window) hold a set of events

An *event handler* can connect to an event

Listens to the stream of data, and performs some action for each item in the stream

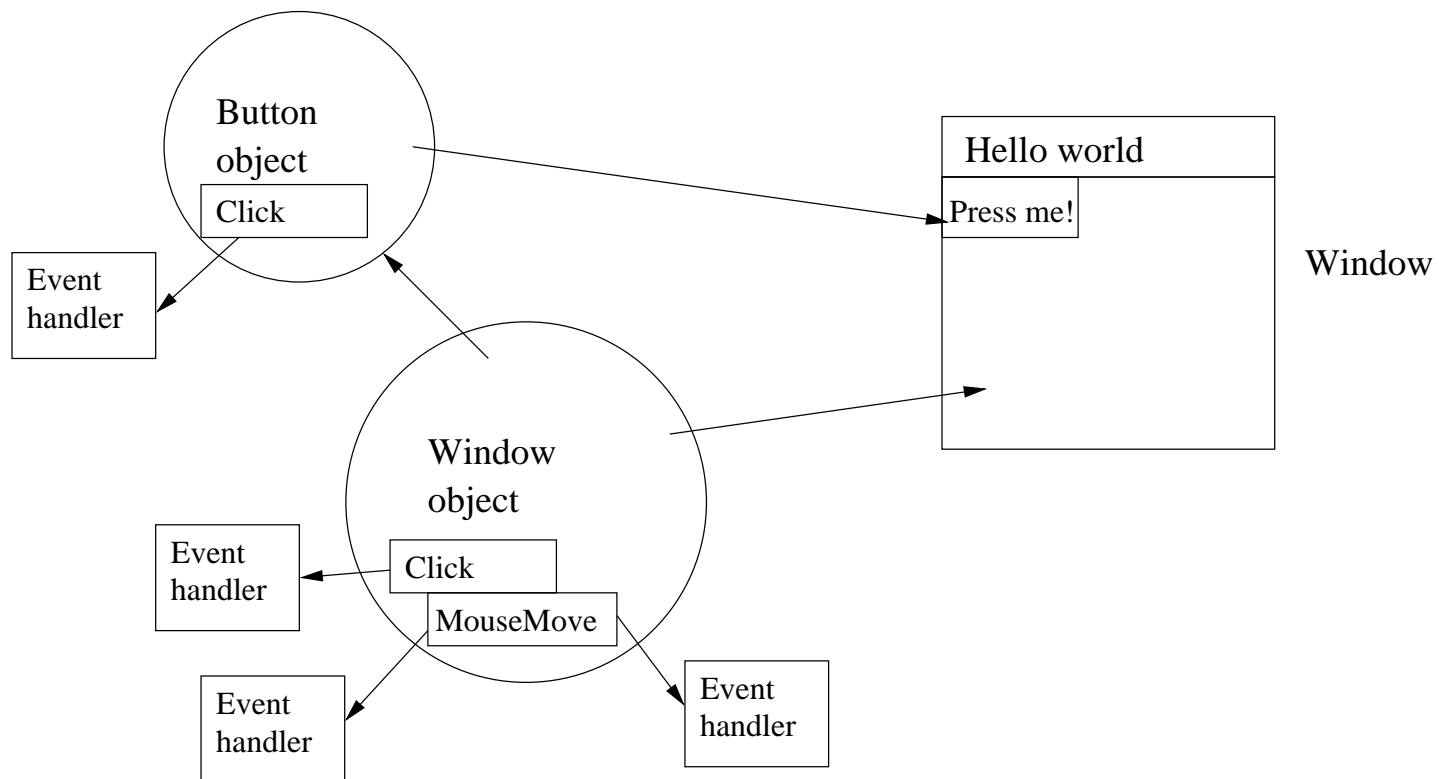


In F#/.NET, an event handler connects to an event by *adding* itself to the event

Once added, it will receive all data in the stream and can take action accordingly

Each GUI object holds a number of events, to which event handlers can add themselves

In F#, event handlers are functions. The function is applied to each data in the stream, in order



Events in F#

F# has a data type `IEvent<'a>` for events

Events are first-class citizens just like any other data: can be moved around, copied, stored in data structures, ...

Since events are data streams, they are similar to sequences

There is an `Event` module with functions on events. Some examples:

```
Event.choose : ('a -> 'b option) -> IEvent<'a> -> IEvent<'b>
Event.filter : ('a -> bool) -> IEvent<'a> -> IEvent<'a>
Event.map : ('a -> 'b) -> IEvent<'a> -> IEvent<'b>
Event.merge : IEvent<'a> -> IEvent<'a> -> IEvent<'a>
Event.partition : ('a -> bool) -> IEvent<'a> ->
    IEvent<'a> * IEvent<'a>
Event.scan : ('a -> 'b -> 'a) -> 'a -> IEvent<'b> ->
    IEvent<'a>
```

Note that some are the same as for sequences (and lists, and arrays)!

The same style of programming can be used for events!

A Simple Example

```
open System.Windows.Forms // Module for .NET GUI handling
open System.Drawing // Namespace for colors etc.

let form = new Form(Text="Hello World",Visible=true)
// Create a new window, and make it visible

let button = new Button(Text="Press here!")
// Create a new button

button.BackColor <- Color.Red
button.Size <- new Size(50,50)
button.Location <- new Point(25,25)
// make it red, size 50x50 pixels, offset (25,25) pixels

button.Click.Add(fun _ -> printfn "You pressed me!!")
// add a handler for the button's "Click" event
```

```
form.Controls.Add(button)
// add the button to the window

form.MouseMove.Add
    (fun args -> printfn "Mouse, (x,y) = (%A,%A)" args.X args.Y)
// add a handler for the window "MouseMove" event

let e = form.MouseMove |> Event.filter (fun args -> args.X > 100)
    |> Event.map (fun args -> args.X + args.Y)
// Define a new event, created from the MouseMove event

e.Add (fun n -> printfn "Mouse sum = %d" n)
// Add an event handler to our new event

Application.Run(form) // Finally start the execution of the window
```