APPLE:
Advanced Procedural Programming Language Elements

Christian Heinlein
Dept. of Computer Structures
University of Ulm
Germany
heinlein@informatik.uni-ulm.de

Position Paper for the ECOOP Workshop on Programming Languages and Operating Systems
(ECOOP-PLOS 2004)

1. Introduction

Today's programming languages, in particular aspect-oriented languages such as AspectJ [9], have received a considerable degree of complexity, making it both hard to learn their "vocabulary" (i. e., simply know all concepts and constructs offered by the language) and to "fluently speak" them (i. e., successfully apply these concepts and constructs in daily programming). In contrast, traditional procedural languages, such as Pascal or C, provided just two basic building blocks: *data structures* (records in particular) and *procedures* operating on them [1]. Modern procedural languages, such as Modula-2 or Ada, added the concept of *modules* to support encapsulation and information hiding [10]. In object-oriented languages such as Eiffel, Java, or C++, these separate and orthogonal entities have been combined into *classes* which offer *subtype polymorphism*, *inheritance* of data structures and procedures (which are usually called *methods* there), and *dynamic binding* of procedures as additional basic concepts.

Even though object-oriented languages support the construction of software that is usually more flexible, extensible, and reusable than traditional "procedural software," it soon turned out that many desirable properties are still missing. For example, *modular extensibility* (i. e., the ability to extend an existing system without modifying or recompiling its source code) is limited to adding new (sub)classes to a class hierarchy, while adding new operations (methods) to existing classes is impossible. Similarly, retroactively extending or modifying the behaviour of operations is infeasible. A great deal of research efforts have been expended in the past years to overcome these limitations by providing even more new concepts, e. g., open classes [3] or advice and inter-type member declarations in aspect-oriented languages [9], to name only a few.

Even though the set of these additional concepts is "sufficient" (in the sense that they indeed solve the problems encountered with object-oriented languages), the question arises whether they are really "necessary" (in the sense that a smaller or simpler set of concepts would not be sufficient). Using AspectJ as an extreme example, this language provides eight more or less different kinds of "procedures," i. e., named blocks of executable code: static methods, instance methods and constructors defined in classes (as in the base language Java), plus instance methods and constructors defined as inter-type members in aspects, plus before, after, and around advice (still neglecting the distinction between "after returning," "after throwing," and general "after" advice).

Figure 1 illustrates this observation graphically: The road leading from procedural languages via object-oriented languages to "conceptually sufficient" aspect-oriented languages climbs up the "hill of complexity" by introducing more and more specialized language constructs in order to "patch" the original deficiencies of procedural and object-oriented languages. This hill of complexity is an undesired burden for language designers and implementors as well as for language users.

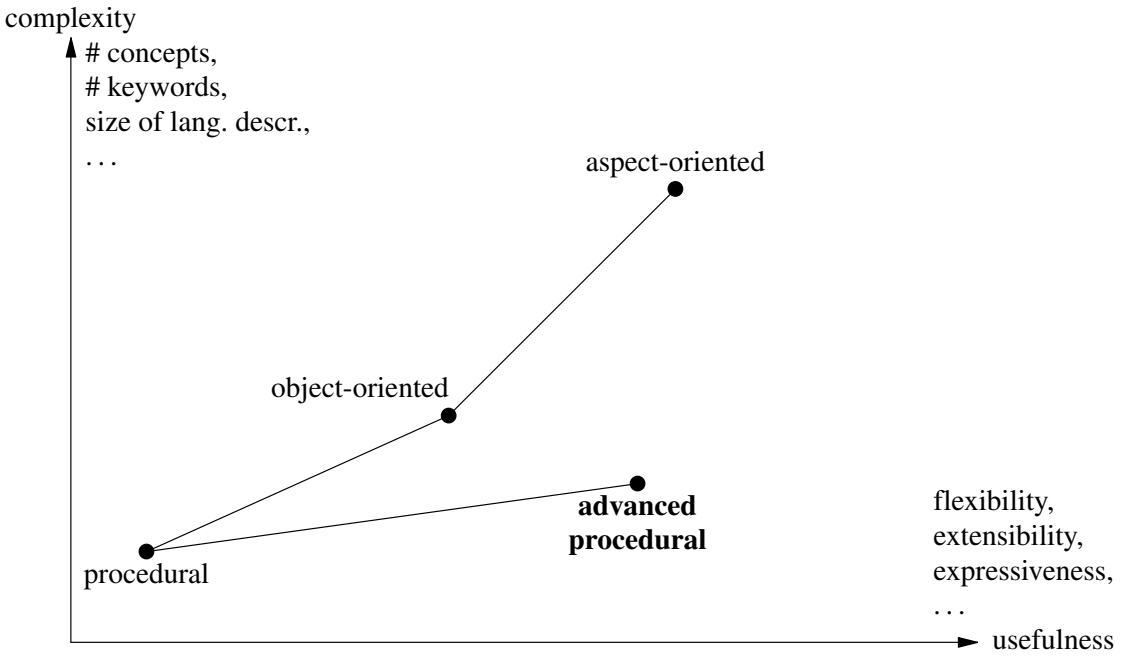

Figure 1: Hill of complexity

From a strictly conceptual point of view, this manifold of procedures is highly redundant: methods and constructors defined in classes are dispensable because they could always be defined in aspects; method and constructor bodies are dispensable because their code could always be defined as advice; before and after advice (the latter in its three variants) are dispensable because they are just special cases of around advice (calling `proceed` at the end resp. beginning of the advice, appropriately embedded in a `try`/`catch` block if necessary to distinguish the three variants of after advice). After these conceptual reductions, around advice – i. e., the possibility to freely override an existing "procedure," either with completely new code (that does not call `proceed`) or with code that augments the original code (by calling the latter via `proceed`) – remains as one of the *essential* (i. e., really necessary) procedure categories. (This goes in line with the statement that "dynamically scoped functions are the essence of AOP" [4].)

It turns out, however, that the potential for conceptual reductions is still not exhausted: By employing dynamic type tests (`instanceof` operator) in an around advice, a programmer is able to emulate the standard dynamic method dispatch provided by the base language Java (or any other dispatch strategy he likes) by executing the advice code only if the dynamic type of the current object is a particular subtype of its static type (or if some other arbitrary condition is satisfied) and simply call `proceed` otherwise. This means in consequence that the specialized built-in dispatch strategy for methods is dispensable from a purely conceptual point of view, thus removing the essential difference between statically and dynamically bound methods, i. e., between static and instance methods.

Similar considerations can be applied to data structures: data fields in classes are dispensable because they are just a special case of inter-type field declarations in aspects. Taken to the extreme, classes can always be declared with empty bodies, because their data fields, constructors, and methods can be declared more modularly and flexibly in aspects.

2. Suggestion

Given these observations, the basic suggestion of this paper is to go back to the starting point of procedural programming languages and extend them into a different direction in order to create *advanced procedural languages* which are significantly simpler than aspect-oriented languages while offering comparable expressiveness and flexibility (cf. Fig. 1).

In particular, replacing simple, statically bound procedures with arbitrarily overridable *dynamic procedures* (roughly comparable to around advice) covers (with some additional syntactic sugar which is not essential) the whole range of dynamic dispatch strategies usually found in object-oriented languages (single, multiple, and even predicate dispatch [2, 5]) plus the additional concept of advice (before, after, and around) introduced by aspect-oriented languages. Nevertheless, dynamic procedures remain a single, well-defined concept which is in no way entangled with data structures, class hierarchies, and the like and therefore is hardly more complex than traditional procedures.

Similarly, replacing simple record types having a fixed set of fields with modularly extensible *open types* and *attributes* (roughly comparable to empty classes extended by inter-type field declarations) covers classes and interfaces, field declarations in classes and aspects, multiple inheritance and subtype polymorphism, plus inter-type parent declarations and advice based on `get` and `set` pointcuts (since reading and writing attributes of open types is implicitly done via overridable dynamic procedures). Again, open types constitute a single, well-defined concept which is little more complex than traditional record types.

Finally, preserving resp. (re-)introducing the *module* concept of modern procedural languages with clearly defined import/export interfaces and a strict separation of module definitions and implementations [12], provides perfect support for encapsulation and information hiding, even for applications where sophisticated concepts such as nested or friend classes are needed in today's languages [6, 11].

3. An Example of Open Types and Dynamic Procedures

This section presents a brief example of open types and dynamic procedures in "Advanced C." A little software library for the representation and evaluation of arithmetic expressions shall be developed.

We start by defining some open types with associated attributes.

```
// General expression.
type Expr;

// Constant expression.
type Const;
conv Const -> Expr; // Const is convertable to Expr, i. e. a subtype.
attr val : Const -> int;        // Value of constant expression.

// Binary expression.
type Binary;
conv Binary -> Expr; // Binary is a subtype of Expr, too.
attr op : Binary -> char;       // Operator and
attr left : Binary -> Expr;     // left and
attr right : Binary -> Expr;    // right operand of binary expression.
```

Then, a dynamic procedure (or *global virtual function* in the nomenclature of C/C++) called `eval` is defined to compute the value of an expression.

```
// Evaluate constant expression.
// The static type of x is Expr, but this "branch" of eval
// is executed only if its dynamic type is Const.
```

```
virtual int eval (Expr x : Const) {
    return x@val; // @ is the attribute access operator
        // similar to the dot operator in other languages.
}

// Evaluate binary expression.
// This branch is executed if x's dynamic type is Binary.
virtual int eval (Expr x : Binary) {
    switch (x@op) {
    case '+': return eval(x@left) + eval(x@right);
    case '-': return eval(x@left) - eval(x@right);
    case '*': return eval(x@left) * eval(x@right);
    case '/': return eval(x@left) / eval(x@right);
    }
}
```

In a later stage of the development, we detect that we have forgotten to implement the remainder operator `%`. We fix this in a completely modular way (i. e., without the need to touch or recompile the above code) by adding another branch of `eval` overriding the previous one if the additional condition `x@op == '%'` is satisfied.

```
// Evaluate remainder expression.
// This branch is executed if x's dynamic type is Binary
// and the condition x@op == '%' holds.
virtual int eval (Expr x : Binary) if (x@op == '%') {
    return eval(x@left) % eval(x@right);
}
```

For a particular application of the library, we might want divisions by zero to return a special null value (represented, e. g., by the smallest available integer value) that propagates through all arithmetic operations (similar to the notion of "not a number" defined by IEEE 754 floating point arithmetics). This can be achieved, again in a completely modular way, by introducing the following additional branches of `eval`.

```
// Special null value.
const int null = INT_MIN;

// Catch divisions by zero.
virtual int eval (Expr x : Binary) if (x@op == '/' || x@op == '%') {
    if (eval(x@right) == 0) return null;
    else return virtual(); // Call previous branch.
}

// Catch null-valued operands.
virtual int eval (Expr x : Binary) {
    if (eval(x@left) == null || eval(x@right) == null) return null;
    else return virtual(); // Call previous branch.
}
```

Note that the order in which the branches are defined is crucial in this example: Since the last branch – which will be tried first when the function is invoked – catches null-valued operands, the second last branch will only be tried if both operands are not null and so does not need to repeat this test.

4. Application to Operating Systems

Even though advanced procedural languages are intended to be general-purpose programming languages, their application to operating systems development might be particularly interesting since many of these systems are still implemented in traditional procedural languages (C in particular). Moving, e. g., from C to an "Advanced C" offering open types and dynamic functions should be much more smooth than shifting to an object-oriented or even aspect-oriented language, since the basic programming paradigm remains the same. Furthermore, by interpreting every standard C function as a dynamic function and every standard C struct as an open type with some initially associated attributes, it is possible to turn existing source code into flexibly extensible code at a glance, by simply recompiling it. With some system-dependent linker tricks it is even possible to turn standard library functions to dynamic functions without even recompiling them.

Operating systems, like software systems in general, usually evolve over time. Taking Unix and its derivatives as a typical example, this system started as a rather small and comprehensible system offering a few basic system calls which implemented a few fundamental concepts. Over the years and decades, it has grown into a large and complex system offering dozens of additional system calls implementing a large number of advanced concepts.

When using conventional programming languages, the introduction of each new concept typically requires modifications to numerous existing functions in addition to implementing new functions. Using open types and dynamic functions instead offers at least the chance to be able to implement new functionality in a truly *modular* way by grouping new functions and necessary redefinitions of existing functions together in a single new unit of code.

To give a concrete example, the introduction of *mandatory file locking* into Unix required modifications to the implementation of several existing system calls (such as `open`, `read`, and `write`) to make them respect *advisory locks* on a file (a concept that has been introduced earlier) as mandatory if the file's access permission bits contain an otherwise meaningless combination. Furthermore, this particular combination of access permissions has to be treated specially at other places of the system, e. g., by not performing the standard action of resetting the "set group ID on execution" bit when such a file is overwritten. By employing dynamic functions, modifications such as these can be implemented without touching or recompiling existing source code by simply overriding existing functions with new functions that perform additional tests before calling their previous implementation or signalling an error such as "lock violation" if necessary.

5. Conclusion

Advanced procedural programming languages have been suggested as an alternative direction to extend traditional procedural languages to make them more flexible and useful. In contrast to object-oriented and aspect-oriented languages, which combine the existing concepts of modules, data structures, and procedures into classes while at the same time introducing numerous additional concepts, advanced procedural languages retain these basic building blocks as orthogonal concepts which are only slightly extended to achieve the primary aim of modular extensibility.

Even though a first version of an "Advanced C" (that is actually being implemented as a language extension for C++ to get for free some of the advanced features of C++, such as templates and overloading of functions and operators) has been used successfully to implement some small to medium-sized programs (and there are also implementations available for dynamic procedures in Oberon and dynamic class methods in Java [7, 8]), it is too early yet to respectably report about experience and evaluation results. Of course, dynamic procedures are less efficient at run time than statically bound procedures because every explicit or implicit delegation of a call to the previous branch of the procedure is effectively another procedure call, at least when implemented straightforwardly without any optimizations. Furthermore, inlining of procedure calls becomes impossible if procedures can be freely re-

defined elsewhere. Nevertheless, the performance penalty encountered appears to be tolerable in practice if the concept is used reasonably.

It is often argued that the possibility to freely redefine procedures anywhere might quickly lead to incomprehensible code because this possibility might indeed be abused to completely change the behaviour of everything in a system. However, the limited practical experience gained so far suggests that the opposite is true, because when applied with care this possibility provides the unique ability to group related code together at a single place instead of needing to disperse it throughout a whole system. By that means, it is possible to develop and understand a system incrementally: Given that the basic functionality of the system is correct, it is possible to reason about its extensions separately in a modular way.

References

[1] A. V. Aho, J. E. Hopcroft: *Data Structures and Algorithms*. Addison-Wesley, Reading, MA, 1983.

[2] C. Chambers, W. Chen: "Efficient Multiple and Predicate Dispatching." In: *Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '1999)* (Denver, CO, November 1999). *ACM SIGPLAN Notices* 34 (10) October 1999, 238–255.

[3] C. Clifton, G. T. Leavens, C. Chambers, T. Millstein: "MultiJava: Modular Open Classes and Symmetric Multiple Dispatch for Java." In: *Proc. 2000 ACM SIGPLAN Conf. on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA '00)* (Minneapolis, MN, October 2000). *ACM SIGPLAN Notices* 35 (10) October 2000, 130–145.

[4] P. Costanza: "Dynamically Scoped Functions as the Essence of AOP." *ACM SIGPLAN Notices* 38 (8) August 2003, 29–36.

[5] M. Ernst, C. Kaplan, C. Chambers: "Predicate Dispatching: A Unified Theory of Dispatch." In: E. Jul (ed.): *ECOOP'98 – Object-Oriented Programming* (12th European Conference; Brussels, Belgium, July 1998; Proceedings). Lecture Notes in Computer Science 1445, Springer-Verlag, Berlin, 1998, 186–211.

[6] J. Gosling, B. Joy, G. Steele: *The Java Language Specification*. Addison-Wesley, Reading, MA, 1996.

[7] C. Heinlein: *Vertical, Horizontal, and Behavioural Extensibility of Software Systems*. Nr. 2003-06, Ulmer Informatik-Berichte, Fakultät für Informatik, Universität Ulm, July 2003.
http://www.informatik.uni-ulm.de/pw/berichte/

[8] C. Heinlein: "Dynamic Class Methods in Java." In: *Net.ObjectDays 2003. Tagungsband* (Erfurt, Germany, September 2003). tranSIT GmbH, Ilmenau, 2003, ISBN 3-9808628-2-8, 215–229.

[9] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, W. G. Griswold: "An Overview of AspectJ." In: J. Lindskov Knudsen (ed.): *ECOOP 2001 – Object-Oriented Programming* (15th European Conference; Budapest, Hungary, June 2001; Proceedings). Lecture Notes in Computer Science 2072, Springer-Verlag, Berlin, 2001, 327–353.

[10] D. L. Parnas: "On the Criteria to Be Used in Decomposing Systems into Modules." *Communications of the ACM* 15 (12) December 1972, 1053–1058.

[11] B. Stroustrup: *The C++ Programming Language* (Special Edition). Addison-Wesley, Reading, MA, 2000.

[12] N. Wirth: *Programming in Modula-2* (Fourth Edition). Springer-Verlag, Berlin, 1988.