Lean and Efficient System Software Product Lines: Where Aspects Beat Objects[★]

Daniel Lohmann, Olaf Spinczyk, and Wolfgang Schröder-Preikschat

Friedrich-Alexander-University Erlangen-Nuremberg, Germany
Computer Science 4 – Distributed Systems and Operating Systems
{dl, os, wosch}@cs.fau.de

Abstract. Software development in the domain of embedded and deeply embedded systems is dominated by cost pressure and extremely limited hardware resources. As a result, modern concepts for separation of concerns and software reuse are widely ignored, as developers worry about the thereby induced memory and performance overhead. Especially object-oriented programming (OOP) is still little in demand. For the development of highly configurable fine-grained system software product lines, however, separation of concerns (SoC) is a crucial property. As the overhead of object-orientation is not acceptable in this domain, we propose aspect-oriented programming (AOP) as an alternative. Compared to OOP, AOP makes it possible to reach similar or even better separation of concerns with significantly smaller memory footprints. In a case study for an embedded system product line the memory costs for SoC could be reduced from 148–236% to 2–10% by using AOP *instead* of OOP.

1 Introduction

The domain of embedded and deeply embedded devices is dominated by 8 bit μ-controllers with 0.25–4 KB of RAM and a few KB of flash memory. In 2000, more than 98% of the total worldwide CPU production (8 billion entities) were dedicated to the domain of embedded systems; 87% of the entities used in this domain were 8 bit or smaller [47]. From the viewpoint of procurement, this "old-fashioned technology" is still the best compromise with respect to functionality and costs. In areas of mass production a *few cents* decide over market success or failure – a situation that cannot be expected to change soon, given that the envisioned scenarios of *smart dust* [30], *ubiquitous computing* [48], and *proactive computing* [47] crucially depend on the bulk availability of very cheap, self-organizing "intelligent" devices.

Counting cents in hardware procurement basically leads to counting bytes in software development. To cope with hardware cost pressure and extremely limited resources, software developers for deeply embedded systems intentionally avoid modern language concepts for a better separation of concerns (SoC), as they worry about the thereby induced memory and performance overhead. Especially object-oriented programming (OOP) is still little in demand, as some of its fundamental concepts (e.g., *late*

---

binding by virtual functions) are known to have non-negligible costs [20]. Hence, most software development for embedded systems is still performed "ad-hoc" in C (often even assembler) with a strong focus on minimizing hardware requirements. Organized reuse and separation of concerns are often considered as less important.

1.1   System Software Product Lines

The common "ad-hoc" application-specific development approach is doomed to fail for reusable system software, such as operating systems or light-weight middleware for embedded devices. System software for this domain has to not only cope with the afore-mentioned resource constraints, but also with a very broad variety of functional and non-functional requirements [10]. It has to be tailorable to provide exactly the functionality required by the intended application, but nothing more. This leads to a *family-based* or *product line* approach, where the variability and commonality among family members is expressed by *feature models* [16]. The tailorability of software product lines depends mostly on the offered level of functional exchangeability (here denoted as *variability*) and functional selectability (here denoted as *granularity*). Both, variability and gran-ularity require a (right-unique) mapping from implementation components to the fea-tures they implement; thus, a good separation of concerns. As the principal overhead of OOP is not acceptable in this domain, we advocate to use *aspect-oriented programming (AOP)* instead.

The AOP, as well as the OOP, provides means for a better separation of concerns. Today, AOP is mostly perceived as an *extension* to OOP, which leads to the impression that it has to induce similar or even higher overheads. This is understandable, as most aspect languages are actually extensions to object-oriented languages like Java. For As-pectJ [31], studies show furthermore that AOP concepts may indeed lead to some extra overhead [21]. Nevertheless, the AOP itself is not limited to OOP. Language extensions have as well been proposed for non-OOP languages (such as AspectC [12]) and multi-paradigm languages (such as AspectC++ [46]). In these languages it is possible to use AOP not only as an extension, but also as an *alternative* to OOP.

1.2   Objectives

This article shows that for the specific requirements of resource-constrained software development, AOP can be superior to OOP. Our focus is *not* on separation of crosscut-ting concerns and other structural benefits of AOP that have earlier been discussed in many papers. Instead, we want to draw attention to a mostly unexplored benefit of AOP: given a well-optimizing static weaver, AOP can provide SoC more *resource-efficiently* than OOP. Our goals are in particular to:

– broaden the perception of AOP. Currently AOP is mostly perceived as an extension to OOP for a better SoC. Understanding it as an *alternative* to OOP with respect to *resource-efficiency* and *hardware costs* is a new point of view.
– attract new communities to AOP. Especially the systems and embedded communi-ties have a profound skepticism regarding the suitability of high-level programming paradigms ("everything beyond C"), as they (have to) favor efficiency over SoC.

Demonstrating that with aspects it can be possible to reach SoC without having to give up efficency would help to broaden the acceptance of AOP.

For these purposes, we conducted a study in which we compare the resource requirements of three implementations of the same embedded software product line. A highly efficient C-based implementation that does not provide SoC, an object-oriented, and an aspect-oriented implementation. We evaluate the OO and AO implementations regarding the techniques used to reach SoC. The focus is on the thereby induced *costs*. While both implementations reach the SoC goals, the OO-based solution leads to significantly higher resource requirements than the AO-based solution, which can compete with plain C.

1.3   Structure of the Paper

The paper is organized as follows: in Sect. 2, we present our case study about an aspect-oriented and an object-oriented implementation of an embedded software product line. Both implementations are analyzed with respect to used idioms and potential cost drivers. In Sect. 4, these implementations are further investigated regarding their resource requirements. The results are discussed in Sect. 5. Section 6 provides an overview of related work. Finally, the paper is briefly summarized in Sect. 7.

2   Scenario

In the following, we present and analyze a case study of an embedded software product line with three different implementations: C-based, OOP-based, and AOP-based. For the sake of comprehensibility, we have conducted this case study in a more application-oriented field. The scenario is an *embedded weather station* program family that should be configurable in various different variants. It is a somewhat typical example of an embedded application software product line. The identified patterns and results are, however, equally relevant for the development of system software, which is our main field of research.

2.1   Overview

A weather station variant basically consists of one or more *sensors* to gain environmental information and one or more *actors* to process the gathered weather data. Figure 1 shows the possible variants of the weather station as a feature model. Note that the list of available sensors as well as actors is expected to grow in future versions, e.g., by additional sensors for wind direction and humidity. Not visible in the feature model, but nevertheless part of the product line definition, is that we have to distinguish between two kinds of actors:

generic actors   are able to process/aggregate information of any set of sensors. *Display* and *XMLProto* are examples of generic actors. The XML representation of weather information, for instance, can easily and automatically be extended for additional sensors.
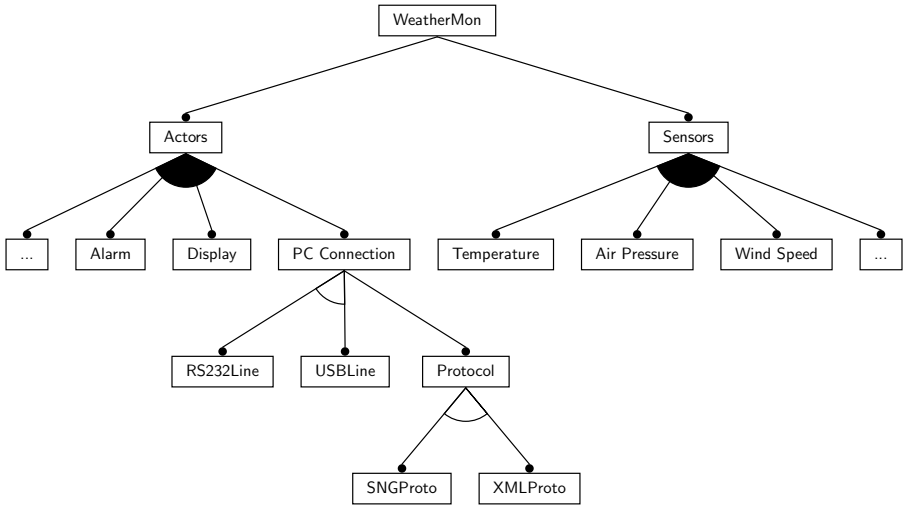
Fig. 1. Feature diagram of the embedded weather station product line. A *WeatherMon* variant consists of at least one sensor and one actor. *Sensors* gain information about the environment, such as *Temperature*, *Air Pressure*, or *Wind Speed*. *Actors* process sensor data: weather information can be printed on a (LCD-) *Display*, monitored to raise an *Alarm* if values exceed some threshold, and passed to a PC through a *PCConnection*, which can be either an *RS232Line* or an *USBLine*, using either an XML-based data representation (*XMLProto*) or a proprietary format (*SNGProto*).

non-generic actors process/aggregate sensor data of some specific sensors only. *SNG-Proto* is an example for a specific actor. It implements a legacy binary data representation for compatibility with existing PC applications. This format encodes information of certain sensors only.

2.2   Hardware Platform

On the hardware side, sensors and actors of the weather station are connected to a small μ-controller (Fig. 2). The AVR series by Atmel is a typical μ-controller product line. It is based on an 8 bit RISC core and offered with a broad variety of on-board equipment: 0–4 KB RAM, 1–128 KB program memory (flash), various busses and connectors (I^2C, serial lines, A/D converters). Wholesale prices of the chip scale between 0.3 and 7 EUR. They depend mostly on the on-board equipment, especially the quantity of RAM and flash memory.

The goal is to reach a similar level of scalability in the software product line, so that hardware costs do scale with the quantity of selected features.

Fig. 2. Weather station hardware platform

2.3   Configuration Process

A weather station variant is configured by selecting features from the feature model. Conceptually, the feature model can be considered as the description of the *configuration space*, while a (valid) feature selection is a description of a concrete variant. In the *implementation space*, a *family model* provides a mapping from the abstract features to concrete implementation artifacts (such as .h and .cpp files) which implement the selected features. The feature selection is evaluated against the configurations space by a variant management tool [9], which copies the thereby defined set of implementation artifacts from the component repository into an output directory. The build process is then started in this directory to compile and link the configured system into the concrete variant (Fig. 3).

The approach to map conceptual features by a family or platform model to concrete implementation artifacts is quite common the domain of embedded system software such as operating system product lines.[1] The main advantages are flexibility and portability. By providing hardware-specific versions of the implementation space (family model and implementation components), the same conceptual OS model and configuration tool can be used to "generate" highly optimized variants for very different μ-controller platforms and derivates.

The overall configurability (variability and granularity) of such product lines depend, however, on the (right-unique) mapping from implementation components to the features they implement; thus, a good separation of concerns. As a matter of fact, many features cannot be mapped to single (C-language) artifacts; their implementation is inter-mingled within the implementation of other features. The result is scattered and tangled code, typically processed by an intra-artifact "second-order configuration" by means of the C preprocessor and conditional compilation. This hampers maintainability and evolvability of the implementation artifacts and thereby limits on the long-term extensibility of the whole product line. Hence, we strive for a better separation of concerns—as long as it does not lead to significant extra costs on the hardware side.

---

[1] Examples are eCos[1], PURE[10], or many implementations of the OSEK system software standard used in the automotive industry [2].
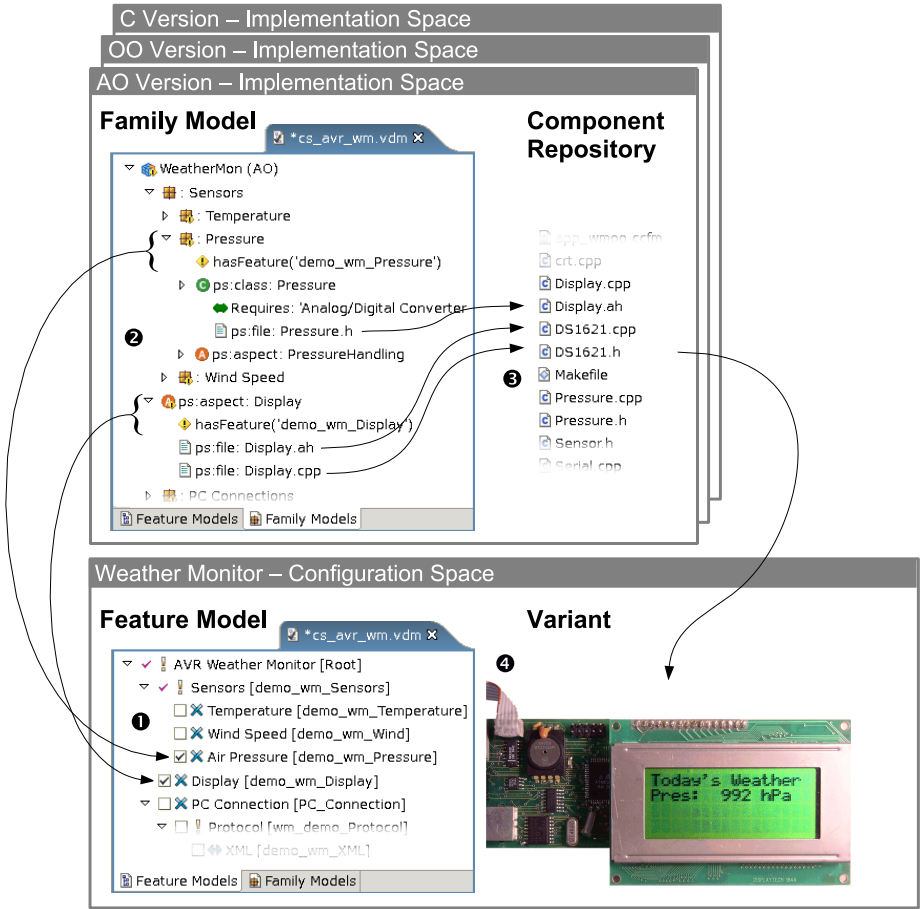
Fig. 3. Weather station configuration process. In the *configuration space* the user selects all wished features from the *feature model* (*1*). The *implementation space* consists of a *family model* and a *component repository*. The family model maps features to *logical implementation components* (*2*), which in turn are mapped to *physical implementation files* (*3*). The thereby determined set of implementation files for a concrete configuration is copied into the target directory, and finally compiled and linked into the actual *weather station variant* (*4*).

3 Implementation

In the following, we present and analyze two different implementations of the weather station product line that fulfill the goal of SoC. One implementation is OOP-based, while the other is AOP-based. The implementations were performed by two different, but equally experienced developers. Both have been (under the premise that SoC has to be achieved) carefully optimized with respect to resource requirements.

A third "traditional" C-based implementation is solely optimized for minimal resource consumption. Its only purpose is to provide the lower bounds of the resource consumption that can be reached in the different product line configurations.

The section is organized as follows: We start with a list of requirements in Sect. 3.1, which is followed by a description of the three implementations (C in Sect. 3.2, OO in Sect. 3.3, and AO in Sect. 3.4). These descriptions are intended as a brief overview only. In particular, they do not motivate the chosen design of the OO and AO versions. Instead, all design decisions and idioms used to reach SoC in the OO and AO versions are discussed collectively (to set them in contrast with each other) in Sect. 3.5. The results are summarized in Sect. 3.6.

3.1   Implementation Requirements

Besides the functional features that have to be implemented, the additional requirements on the implementations can be summarized as *resource-thriftiness* (all versions) and *separation of concerns* (AO and OO versions). This means in particular:

granularity. Components should be fine-grained. Each implementation element should be either mandatory (such as the application main loop) or dedicated to a single feature only.

economy. The use of expensive language features should be avoided as far as possible. For instance, a method should only be declared as `virtual` if polymorphic behavior of this particular method is required.

pluggability. Changing the set of selected sensors and/or actors should not require modifications of any other part of the implementation. This basically means that sensor/actor implementations should be able to integrate themselves into the system.

extensibility. The same should hold for new sensor or actor types, which may be available in a future version of the product line.

3.2   Implementation of the C Version

Figure 4 shows the basic structure of the C version. Weather information is updated by the `measure()` function, which invokes—configured by conditional compilation—all existing sensors. For efficiency reasons, the actual weather data are passed as a global variable (not shown). Weather information is processed by the `process()` function, which invokes each actor (such as `display_process()`)—again configured by means of conditional compilation. The `..._process()` function of each actor retrieves the weather information for all configured sensors it is interested in—once more using conditional compilation.

Due to the design goal to minimize the resource consumption, the C version does *not* fulfill our requirements on SoC (*granularity*, *pluggability*, *extensability*). By using global variables, conditional compilation, and inlining of all functions that are referred only once, it offers, however, efficiency.

init_sensors (main.c)

```
inline void init_sensors(){
#ifdef cfWM_WIND
  wind_init();

#endif
#ifdef cfWM_PRESSURE
  pressure_init();

#endif
  ...
}
```

measure (main.c)

```
inline void measure(){
#ifdef cfWM_WIND
   wind_measure();

#endif
#ifdef cfWM_PRESSURE
   pressure_measure();

#endif
   ...
}
```

display_process (display.h)

```
inline void display_process(){
  char val[ 5 ]; UInt8 line = 1;
#ifdef cfWM_WIND
   wind_stringval( val );
   display_print(line++, val,...);
#endif
#ifdef cfWM_PRESSURE
   pressure_stringval( val );
   display_print(line++, val,...);
#endif
   ...
}
```

• • •

(left margin, rotated: *sensor integration*)

init_sinks (main.c)

```
inline void init_sinks(){
#ifdef cfWM_DISPLAY
   display_init();
#endif
#ifdef cfWM_PCCON_XML
   XMLCon_init();
#endif
   ...
}
```

process (main.c)

```
inline void process(){
#ifdef cfWM_DISPLAY
   display_process();
#endif
#ifdef cfWM_PCCON_XML
   XMLCon_process();
#endif
   ...
}
```

main (main.c)

```
int main() {
   ...
   init_sensors();
   init_sinks();
   asm("sei");
   while(true) {
     measure();
     process();
     wait();
   } }
```

(left margin, rotated: *actor integration*)
(margin, rotated: *main loop*)

Fig. 4. Static structure of / scattered code in the C version (excerpt)

3.3   Implementation of the OO Version

Figure 5 shows the class model of the OO version. Central elements are the `Weather` and `Sink` classes. `Weather` aggregates all sensors, which are registered at run time by calling `Weather::registerSensor()`. The `Sink` class aggregates all actors, respectively. Internally both, sensors and actors, are managed by chaining them into light-weight single-linked lists (`ChainBase`).

Principle of Operation

1. Weather information is acquired by calling `Weather::measure()`, which in turn invokes the `Sensor::measure()` method on every registered sensor to update the sensor data.
2. Weather information is processed by the `Sink::process()` method. `Sink::process()` first calls `Actor::before_process()` for each registered actor to initialize the processing.
3. Sensor information is passed to the actors by calling `Actor::process()` for each registered sensor. Actors retrieve the actual sensor name, unit, and measured data (as character strings) by calling the respective Sensor methods.
4. Finally, data processing is finalized by `Sink::process()` invoking `Actor::after_process()` on each actor.
5. The whole process is repeated by the application main loop every second.

3.4   Implementation of the AO Version

The class/aspect model of the AO version is shown in Fig. 6. Central elements are, again, the classes `Weather` and `Sink`. Each sensor class is accompanied by a *han-*
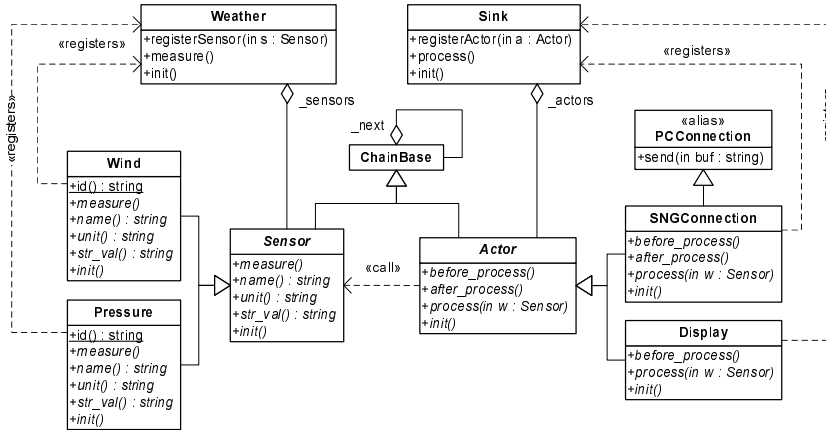
Fig. 5. Static structure of the OO version (excerpt with two sensors/actors, virtual functions are depicted in *italics*, static functions are underlined)

dling aspect, which performs the actual integration into `Weather` and `Sink`. A handling aspect performs two introductions: it aggregates the sensor as an instance variable into class `Weather` and a sensor-specific (empty) `process_data()` method into class `Sink`. Additionally, it defines function execution advice for `Weather::measure()` and `Sink::process`. Actors are implemented as aspects, which define execution advice for `Sink::process()` (for initialization and finalization) and the sensor-introduced `Sink::process_data()` methods (for the actual data processing).

Principle of Operation

1. Weather information is acquired by calling `Weather::measure()`, which is advised by the handling aspects of every sensor to call the sensor's `measure()` method.
2. Weather information is processed by the `Sink::process()` method. `Sink::process()` is before-advised by any actor that needs to initialize before processing.
3. `Sink::process()` is advised for each sensor to call the introduced sensor-specific `process_data()` method.
4. The sensor-specific `process_data()` method is advised by every actor that processes data of this sensor.
5. Finally, data processing is finalized by `Sink::process()` being after-advised for any actor that needs to finalize its processing.
6. The whole process is repeated by the application main loop every second.

3.5   Used AOP and OOP Idioms

To achieve the required level of SoC as well as pluggability and extensibility, both versions use approach-specific idioms and patterns in the design and implementation. In
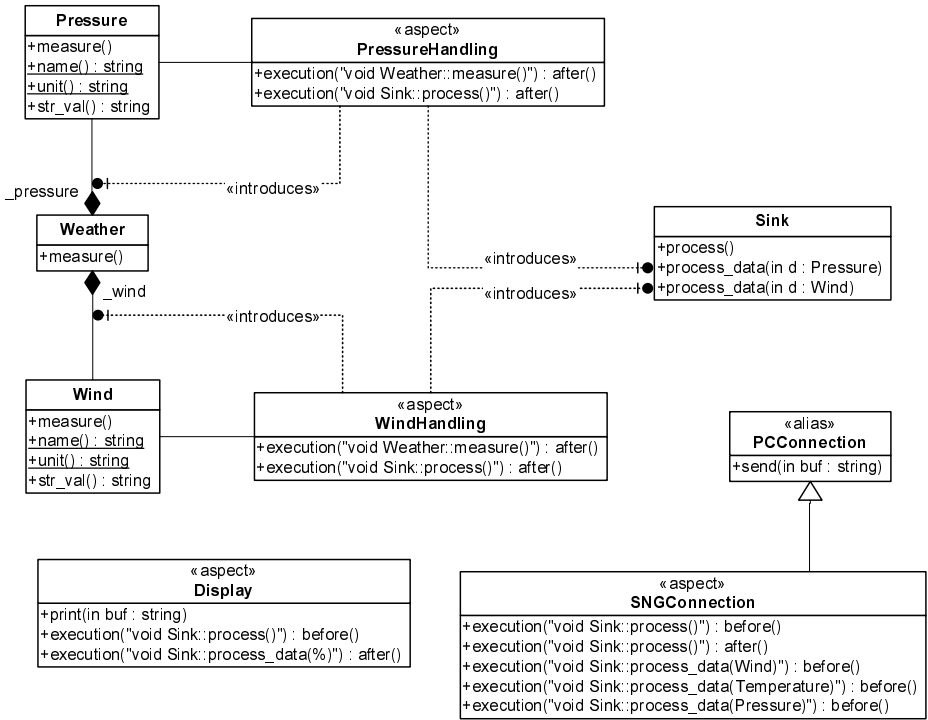
Fig. 6. Static structure of the AOP version (excerpt with two sensors/actors)

the following, we will analyze some interesting parts of the AOP and OOP implementations. The goal is to identify, compare, and discuss the idioms that have to be used to reach SoC. In particular, the following problems had to be solved.

3.5.1 Working with Configuration-Dependent Sensors/Actors Sets

Both implementations have in common that the `Weather`/`Sink` classes are used to abstract from the configured sets of sensors and actors. These two abstractions are unavoidable, because the `main()` function, which performs the endless measurement and processing loop, should be configuration-independent and, thus, robust with respect to extensions. However, the implementation of the 1:n relationship of `Weather`/`Sink` and the concrete sensors and actors is completely different.

Interface dependency between Weather/Sink and sensors/actors
The `Weather`/`Sink` classes need to be able to invoke the measuring/processing of data independent of the actual sensor/actor *types*. Otherwise, pluggability and extensibility would be broken:

OO version: In the OO version, this is solved by *common interfaces* and *late binding*. Sensors have to inherit from the abstract class `Sensor`, actors from the

abstract class `Actor`, respectively. `Weather`/`Sink` invoke sensors/actors via these interfaces, and thus depend on them. `Sensor::measure()`, `Actor::before_process()`, `Actor::after_process()`, and `Actor::process()` have to be declared as virtual functions:

```
struct Sensor : public ChainBase {
  virtual void measure() = 0;
  ...
};
...
class Weather {
public:
  ...
  void measure () {
    for( Sensor* s = ... )
      s->measure();        // virtual function call!
  }
};
```

AO version: In the AO version, the interface relationship is reverted. `Weather` and `Sink` do not depend on sensors/actors providing any specific interface. Sensors/actors are integrated by defining function execution advice for `Weather::measure()` and `Actor::process()`:

```
class Weather {
public:
  void measure () {} // empty implementation
};
...
aspect PressureHandling {
public:
  advice "Weather" : Pressure _pressure;
  advice execution ("void Weather::measure()") : before () {
    // non-virtual inlineable function call
    theWeather._pressure.measure ();
  }
  ...
};
```

Potential cost drivers: In the OO version, four methods in two classes have to be declared and called as virtual functions. The AO version induces no costs, as all sensor/actor code can be inlined.

Working with sets of sensors/actors

`Weather`/`Sink` need to be independent of the actual *number* of sensors/actors configured into the system. There is a 1:n relationship between `Weather`/`Sink` and sensors/actors:

OO version: In the OO version, this is solved by a simple *publisher/subscriber* mechanism. Sensors are chained into a linked list of "publishers", which is frequently iterated by `Weather::measure()` to update the data of each sensor:

```
class Weather {
  static  Sensor* _sensors;
public:
  Sensor* firstSensor() const {
    return _sensors;
  }
  void measure () {
    for( Sensor* s = firstSensor(); s != 0;
         s = static_cast< Sensor* >( s->getNext() ) )
      s->measure();
```

```
        }
        ...
    };
```

Actors are similarly chained into a list of "subscribers" which are invoked if new sensor data are available. The `Sink` class acts as "mediator" between sensors and actors. For the sake of memory efficiency, actors do not subscribe for single sensors, but are implicitly subscribed for the complete list.

AO version: In the AO version, actors and sensors are implicitly chained at compile time by multiple aspects defining advice for the same `Weather::measure()` and `Sink::process()` execution joinpoints.

Potential cost drivers: In the OO version, `Weather`/`Sink` as well as sensors/actors need to carry an extra pointer for the chaining. Some additional code is required to iterate over the chain of sensors/actors.[2]

Registration of sensors/actors
This problem is closely related to the previous one. Sensors and actors need to be able to register themselves for the publisher/subscriber chain:

OO version: In the OO version, registration is done at run time by calling `Weather::registerSensor()`/`Sink::registerActor()`. Self-registration requires some C++ trickery: Each sensor/actor is instantiated as global object. The registration is performed by the constructor, which is "automatically" triggered during system startup.

AO version: In the AO version, no extra efforts are required for self-registration. The chaining of actors and sensors is implicitly performed at compile time by advice-code weaving. Basically, it is the presence of some aspect in the source tree that triggers the registration.

Potential cost drivers: In the OO version, the use of global object instances causes some overhead, as the compiler has to generate extra initialization code and puts a reference to this code in a special linker section. A system startup function has to be provided that iterates over this section and invokes the constructors of all global instances.

3.5.2 Implementation of Generic Actors
Generic actors process sensor data from any sensor type. For example, the Display actor should print the measured values of any configured sensor, regardless of the current configuration or future extensions with new sensor types. However, this leads to an interface dependency between actors and sensors, because (at least) the value has to be obtained in a generic way:

---

[2] The `firstSensor()` and `getNext()` operations are, however, inlined as they perform just a pointer lookup.

OO version: In the OO version, this is again solved by *interfaces* and *late binding*. To enable actors to retrieve sensor information from any sensor, the `Sensor` interface has to be extended by three additional virtual functions: `Sensor::name()`, `Sensor::unit()`, and `Sensor::str_val()`:

```
class Display : public Actor {
public:
    UInt8 _line;

  // print a line on the display, increment _line
  void print(const char *name, const char *val_str, const char *unit );
  ...
  virtual void before_process() {
    _line = 1;
  }
  // called by Sink::process for every sensor
  virtual void process( Sensor* s ) {
    char val[ 5 ];
    s->str_val( val );
    print( s->name(), val, s->unit() );
  }
};
```

AO version: In the AO version, this is solved by using the AspectC++ concept of *generic advice* [32, 35]. For every sensor `MySensor`, the corresponding `MySensorHandling` aspect introduces an empty sensor-specific `process_data(const MySensor&)` method into class `Sink` and gives advice to `Sink::process()` to invoke the introduced method:

```
aspect PressureHandling {
public:
  advice "Weather" : Pressure _pressure;
  ...
  // introduce an empty process_data function for the pressure
  advice "Sink" : void process_data( const Pressure & ) {}
  // call Sink::process_data for the pressure
  advice execution("void Sink::process()") : after () {
    theSink.process_data( theWeather._pressure );
  }
};
```

An actor gives generic advice that matches every (overloaded) `Sink::process_data()` function. To be independent of the actual sensor type, the advice body uses the AspectC++ joinpoint API to retrieve a typed reference to the sensor instance:

```
aspect Display {
  ...
  // display each element of the weather data
  advice execution("void Sink::process_data(%)") : before () {
    typedef JoinPoint::template Arg<0>::ReferredType Data;
    char val[5];
    tjp->arg<0>()->str_val( val );
    print( Data::name(), val, Data::unit() );
  }
};
```

AspectC++ instantiates advice bodies per joinpoint. Therefore, the calls to the actual sensor's `str_val()`, `name()`, and `unit()` methods can be bound at compile time. As an additional optimization, `name()` and `unit()` are implemented as static (class) functions.

Potential cost drivers: In the OO version, three additional virtual functions are required in the sensor classes, as well as an additional virtual function calls in the actor classes.

In the AO version, the joinpoint API has to be used for a uniform and type-safe access to the sensor instance, which induces some overhead. Furthermore, the advice body of a generic actor is instantiated once per sensor, which may lead to code bloating effects.

3.5.3   Implementation of Non-generic Actors

Non-generic actors process data of some sensors only. The legacy SNG protocol, for instance, encodes weather data in a record of wind speed, temperature, and air pressure. It exposes the actual data using sensor-specific interfaces. The record may be sparse, meaning that a specific sensor, such as Temperature, may or may not be present in the actual system configuration.

OO version: In the OO version, a non-generic actor filters the sensors it is interested in by applying *run time type checks* in the `process()` method. Each passed sensor is tested against the handled sensor types. If the run-time type matches with a handled sensor, a downcast is performed to get access to the sensor-specific interface:

```
class Pressure : public Sensor {
  static const char *_id;
public:
  ...
  static const char * id () { return _id; }
  const char *name () const { return _id; }
};
...
class SNGConnection : public Actor, protected PCConnection {
    UInt8 _p, _w, _t1, _t2;     // weather record
public:
  virtual void before_process() { ... /* init record */ };
  virtual void after_process() { ... /* transmit record */ };
  ...
 // collect wind, pressure, temperature data
  virtual void process( Sensor* s ) {
  const char *name = s->name();
    if (name == Wind::id()) // pointer comparison
      _w = ((Wind*)s)->_w;
    else if (name == Pressure::id())
      _p = ((Pressure*)s)->_p - 850;
    else if (name == Temperature::id()) {
      _t1 = (UInt8)((Temperature*)s)->_t1;
      _t2 = ((Temperature*)s)->_t2;
  } }
};
```

The idiom commonly used for run-time type checks in C++ is the `dynamic_cast` operator, which is part of the C++ *run-time type interface* (*RTTI*). RTTI is, however, quite expensive, as it requires additional run-time support and leads to some extra overhead in *every* class that contains virtual functions. To avoid this overhead, our implementation uses a "home-grown" dynamic type-check mechanism: The test is performed by comparing the string address returned by the (late-bound) `Sensor::name()` method with the address of the name string stored in the concrete class,[3] which is also returned by the static `Sensor::id()` method. The expensive C++ RTTI mechanism has been disabled.

Normally, dynamic type checks are considered harmful, because of a lack of extensibility and the accumulated costs of type checks, which sometimes outweigh the costs

---

[3] This basically reduces the overhead of a run-time type test to a virtual function call and a pointer comparison. As the storage for the name string has to be provided anyway, this mechanism also induces no extra overhead in the sensor classes.

of a single virtual function call. However, in our case a *non-generic* actor shall be implemented. Therefore, extensibility is not an issue and the overhead of our type check implementation is acceptable. At the same time, alternative designs such as a *visitor* [24] fail, because a Visitor interface would have to list a visit*Sensor*() method for every sensor type. However, the set of sensors is configurable and should nowhere be hard-wired.

AO version: In the AO version, binding an actor to selected sensors only is realized by giving advice for specific Sink::process_data() methods only instead of using generic advice:

```
aspect SNGConnection : protected PCConnection {
  UInt8 _p, _w, _t1, _t2;      // weather record
  ...
  // let this aspect take a higher precedence than <Sensor>Handling
  advice process () : order ("SNGConnection", "%Handling");

  advice execution("void Sink::process(const Weather&)")
      : before () { ... /* init record */ }
  advice execution("void Sink::process(const Weather&)")
      : after () { ... /* transmit record */ }

  // collect wind, pressure, temperature data by giving specific advice
  advice execution("void Sink::process_data(...)") && args (wind)
      : before (const Wind &wind) {
    _w = wind._w;
  }
  advice execution("void Sink::process_data(...)") && args (pressure)
      : before (const Pressure &pressure) {
    _p = pressure._p - 850;
  }
  advice execution("void Sink::process_data(...)") && args (temp)
      : before (const Temperature &temp) {
    _t1 = (UInt8)temp._t1;
    _t2 = temp._t2;
  }
};
```

Potential cost drivers: Run-time type checks in the OO version induce nevertheless some overhead. In the AO version, some overhead is induced by the args() pointcut function, which is used here to get the actual sensor instance.

3.6 Implementation Summary

Both, the OOP as well as the AOP implementation of the embedded weather station product line provide good SoC. In particular, the implementation requirements described in Sect. 3.1 are met by both versions:

granularity is achieved by the OO as well as the AO version. Each implementation component is either mandatory (such as the Weather and Sink classes), or dedicated to a single feature only.

economy is achieved as far as possible. In the OO version, only methods that have to be available via a generic interface are declared as virtual. RTTI is not used, as the required run-time type checks can be implemented with less overhead. In the AO version, joinpoint-specific context information is used only sparingly.

pluggability is achieved as well. In both versions, no component has to be adapted if the set of selected sensors/actors is changed. Sensors and actors basically integrate themselves, if their implementation component is present in the configured source tree. The OO version uses global instance construction for this purpose. In the AO version, the integration is performed by advice.

extensibility is also achieved. In the OO version, new sensor/actor types just need to implement the common `Sensor`/`Actor` interface. In the AO version, new sensor/actor types just need to provide some aspect that performs the integration into `Weather`/`Sink`.

Overall, the AO and OO versions are equipollent from the SoC viewpoint. We, however, identified noticeably more potential cost drivers in the OO version than in the AO version. Especially virtual functions were unavoidable in many places to realize loose coupling and genericity of components. In the next section, we analyze how this affects scalability and memory demands of the product line.

4   Cost Analysis

In this section, we analyze scalability and memory requirements of the embedded weather station product line. For this purpose, several configurations of the weather station were generated as AO, OO, and C variants.[4] For each variant, we measured:

- static memory demands, which are determined by the amount of generated machine code (*text*), static initialized data (*data*), and static non-initialized data (*bss*).
- dynamic memory demands, which are determined by the maximum stack space used by the running application (*stack*).[5]
- the run time of a complete measure/process-cycle.

On the actual hardware, *text* occupies flash memory space. *Data* occupies flash memory *and* RAM space, as it is writable at run time and therefore has to be copied from flash into RAM during system startup. *Bss* and *stack* occupy RAM space only.

4.1   Measurement Methods

Static memory demands (*text*, *data*, *bss*) could easily be retrieved directly from the linker map file. Dynamic memory demands (*stack*) and run time had to be measured in the running targets:

---

[4] All variants were compiled with `avr-g++` (GCC) 3.4.1 using `-Wall -fno-rtti -Os -fno-exceptions -fomit-frame-pointer -ffunction-sections` optimization flags. AO variants were woven with `ac++` 0.9.3.

[5] The weather station software uses no heap, which otherwise would also contribute to dynamic memory demands.

4.1.1 Measuring Stack Utilization

For simple programs it is possible to determine the stack utilization offline and byte-exact by static analysis of the machine code. Unfortunately in our case the program execution graph is not predictable due to the use of interrupts and late bound functions. Therefore we used run-time monitoring of the stack as a pragmatic alternative. A common technique for run-time stack monitoring is to initialize the entire stack space with some specific *magic pattern* during system startup. The maximum quantity of stack used can then be measured at any time by searching (from bottom of stack) for the first byte where the pattern has been overwritten.

In the weather station variants, this technique was used to implement stack measurement as an additional *sensor type*. Understanding stack measurement as just another sensor had some nice advantages. Because of the achieved pluggability (in the AO and OO versions) it was very easy to apply stack measurement to any weather station configuration. Of course, some extra care had to be taken to ensure that the maximum stack utilization is not caused by the stack measurement sensor itself. For this reason, the stack measurement implementation uses only global variables which do not occupy stack space. By analyzing the generated machine code we ensured that the stack utilization of the stack sensor methods is minimal among all sensors. As all sensor methods are invoked from the same call depth level, and at least one "real" sensor besides stack measurement is used in a weather station configuration, it can thereby be guaranteed that stack measurement itself does not tamper the maximum stack utilization. In the actual targets, the thereby acquired maximum stack utilization remained stable after a short startup time and could be read from one of the attached generic actors.

4.1.2 Measuring Run Time

Run-time measuring was not implemented as another sensor type, as it had been too difficult to distinguish the run time taken by the sensor processing itself from the run time of the target to measure. Instead, we used a less invasive approach that could be implemented with just two additional assembler statements: in the application main loop, a digital I/O port of the AVR μ-controller is set to *high* before the call to `Weather::measure()` and reset to *low* after the return of `Sink::process()`. The result is a rectangular signal on this port, which was recorded and analyzed with a storage oscilloscope.[6] A *high* period in the signal represents the run time taken by a complete `Weather::measure()` / `Sink::process()` cycle. After a short startup time, period and phase of the signal remained stable, and the length of the *high* phase could be measured.

4.2 Overall Scalability of the Product Line

As the graphs in Fig. 7 show, the resulting RAM/flash demands of the weather station software do scale quite well with the amount of selected features. The "Barometer" configuration (P+Display), consisting of just an air pressure sensor and an LCD

[6] Tektronix TDS 2012, 100 MHz resolution.

Fig. 7. Footprint comparison (RAM, flash) for different configurations

display, induces significantly smaller memory demands than the "Deluxe-PC" configuration (TWP+Serial+XML+Display) which bundles three sensors, a LCD display, and an XML-based PC connection over a serial interface. The memory requirements of the other examined configurations are in-between. The noticeably high amount of flash memory required by the "Thermometer" configuration (T+Display) can be explained by the fact that this sensor is connected via the I^2C bus to the μ-controller (see Fig. 2). To drive this bus, additional driver code is required that has not been included for other sensors.

Overall, all three versions meet the goal of scalability, which is an indicator for achieved granularity. In every case, however, the OOP version requires significantly more memory space than its AOP counterpart which comes very close to the C version. Depending on the configuration, the required quantity of RAM is up to 138% higher in the OO version, while the AO version takes only an extra of 10% at maximum (up to 13 byte)—both compared to the C-based version that does not provide SoC. The amount of flash memory is up to 91% higher in the OO version, but only 4% at maximum in the AO version. The net difference between using AOP and OOP has to be considered as even higher, as both versions of each configuration are linked with the same
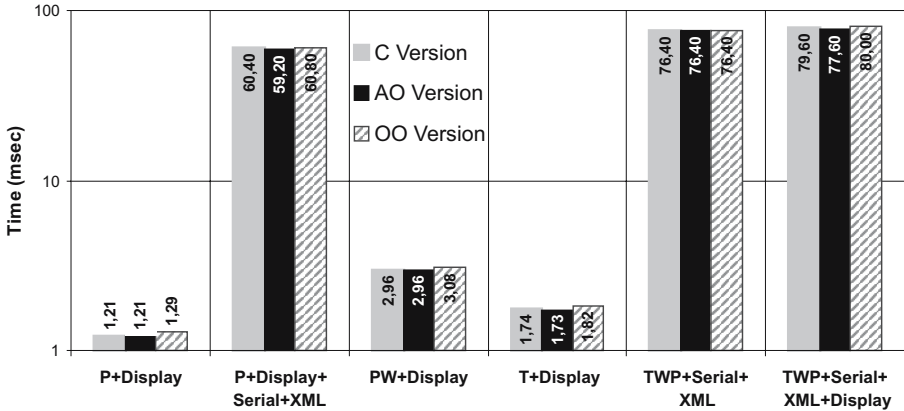
Fig. 8. Footprint comparison (time taken by a `Weather::measure()` / `Sink::process()` cycle) for different configurations

(configuration-dependent) set of device drivers whose memory requirements are included in these numbers.[7]

The run time of all three versions is almost the same in all configurations (see Fig. 8). It is mostly dominated by hardware and driver-related costs.

4.3 Memory Requirements in Detail

Table 1 breaks down the required overall amount of RAM and flash memory into their origins. It is evident, that the OO variants induce especially higher static memory demands. The *text* sections of the OO variants are up to 78%, the *data* sections up to 284% bigger than in the AO variants. The following cost drivers can be mainly accounted for this:

virtual functions are the main source of additional code and data, as they induce overhead on both, the caller and the callee side. On the caller side, each call to a virtual function requires, compared to advice code inlining, at least 16 additional byte in the machine code. On the callee side, virtual function tables have to be provided (4 byte + 2 byte per entry, *data*), object instances need to carry a pointer to the virtual function table (2 byte per instance), and constructors to initialize the vtable pointer of an instance have to be generated (at least 24 byte per class, *text*). In the "Barometer" configuration (P+Display), for instance, 52 byte of the *data* section are occupied solely by virtual function tables.[8]

Regarding code size, the situation may become even worse in larger projects: due to late binding, the bodies of virtual functions are never removed from the final image

---

[7] As parts of the driver code are implemented as inline functions, it is not possible to differentiate between driver-induced and application-induced code here.

[8] The AVR RISC core uses a Harvard architecture; thus vtables cannot be placed in the text section.

Table 1. Memory usage and run time of the AO and OO versions for different configurations

Configuration		text	data	bss	stack	flash	ram	time
P+Display	C	1392	30	7	34	1422	71	1.21
	AO	1430	30	10	38	1460	78	1.21
	OO	2460	100	22	44	2560	166	1.29
P+Display+	C	1578	104	7	34	1682	145	60.40
Serial+XML	AO	1622	104	12	38	1726	154	59.20
	OO	3008	206	26	44	3214	276	60.80
PW+Display	C	1686	38	14	55	1724	107	2.96
	AO	1748	38	18	61	1786	117	2.96
	OO	3020	146	33	65	3166	244	3.08
T+Display	C	2378	28	8	34	2406	70	1.74
	AO	2416	28	11	38	2444	77	1.73
	OO	3464	98	23	44	3562	165	1.82
TWP+Serial+	C	2804	90	17	35	2894	142	76.40
XML	AO	2858	90	23	41	2948	154	76.40
	OO	4388	248	39	41	4636	328	76.40
TWP+Serial+	C	3148	122	17	57	3270	196	79.60
XML+Display	AO	3262	122	24	63	3384	209	77.60
	OO	5008	300	44	67	5308	411	80.00

flash := text + data   ram := bss + stack + data   (Bytes)   (ms)

by means of function-level linking. Thus, "dead" function code that is never called at run time becomes nevertheless part of the *text* section.

dynamic data structures are another source of additional overhead. The chaining of actors and sensors induces 8 additional data byte in the "Barometer" configuration, plus some extra code to access and iterate over the lists.

global instance construction causes some more "hidden" code to be generated. For each translation unit that defines one or more global objects, the compiler has to generate a specific initialization-and-destruction function (88 bytes, *text*). Pointers to these functions are additionally stored in the *data* section.

Regarding dynamic memory usage (*stack*), the differences between the AO and OO version are less significant. The OO variants need a few byte (up to 16%) more stack space than the related AO variants. This seems surprising at first, given that virtual function calls cannot be inlined, which therefore lead to a higher call depth and, thus, higher stack utilization. Part of this effect can be explained by the fact that the AOP version requires some additional stack space as well, namely by context-binding pointcut functions and the joinpoint API (2–4 byte). The main reason is, however, that the maximum virtual function call depth is with two levels quite low. As the AVR architecture provides 32 general-purpose registers, which are also used for passing function parameters, it can furthermore be considered as quite "stack-friendly". On other CPU architectures (such as Intel), the differences between AO- and OO-based solutions would be more significant.

4.4   Run-Time Requirements in Detail

Table 1 also lists the measured run times in detail. In all configurations that support a serial connection the measurement/processing cycle time is mainly dominated by the

underlying serial device driver. Here, the costs are almost the same. In all other configurations the performance of the AO version and our highly efficient C implementation are the same. The run-time overhead of the OO version is between 4 and 6.6%. It can be explained with the numerous virtual function calls on the application level.

4.5   Cost Analysis Summary

Both, the AO and OO version of our product line, scale quite well. The run-time differences are small and only of minor importance in this domain. The OO version, however, induces dramatically higher memory requirements. Given a hardware product line like the AVR ATmega series, these differences can directly be mapped to the required μ-controller features. As Fig. 9 shows, the hardware costs of the C-based and AO-based product line would be exactly the same in all configurations. Using AO versus OO leads even for our small example product line to significant differences regarding hardware costs. This effect would probably be even higher for larger product lines.


Fig. 9. Hardware scalability (required μ-controller variant) of the AO and OO versions for different configurations (wholesale prices: Digi-Key Corporation, Product Catalogue, Summer 2005)

5   Discussion

As pointed out in the previous sections, the AOP-based version of the weather station product line combines good SoC with high memory efficiency. In the following we discuss the principal reasons for this effect, the potential disadvantages, and to what degree our results are applicable to (system software) product lines in general and achievable with other AOP approaches.

5.1   General Advantages of the Approach

To ensure variability and granularity in product lines, components need to be implemented in a fine-grained and loosely coupled way. In our case study, we identified three

fundamental concepts that are required to reach this: *Abstraction from concrete components* (interfaces), *abstraction from component sets* (1:n relationships), and *abstraction from component integration* (self-pluggability). The major disadvantage of OOP is that the features it provides to reach these goals are basically run-time concepts, while with AOP (in the case of static weaving) it is possible to use compile-time concepts:

abstraction from concrete components requires to use interfaces and run-time binding in OOP (thus, virtual functions in C++), while in AOP advice code is woven at compile time to the joinpoints (and in AspectC++ is even inlined).

abstraction from component sets requires to use some kind of run-time data structures in OOP, such as linked lists, while in AOP multiple advice can be woven at the same joinpoint. Moreover, by means of aspect ordering, it is even possible to define a context-dependent invocation order.

abstraction from component integration is generally difficult to achieve in most OOP languages. In C++ it requires some constructor trickery to perform the integration at run time. In AOP, Advice is *per se* an integration mechanism and resolved at compile time. Moreover, aspects provide means for context-sensitive integration: Giving advice for a specific joinpoint can be understood as a *weak reference* to another component, as the advice is silently ignored if the joinpoint does not exist.

5.2   Compiler Issues

As pointed out in the analysis, much of the extra overhead of the OO version is directly or indirectly induced by the costs of late binding, even though the extra flexibility at run time is not required in every case. Much work has been conducted to reduce the costs of C++ virtual functions [3, 8] and late binding in general, e.g., by static whole program optimization [18]. This gives rise to the question, if and to what extent the measured OO overhead could be prevented by better compilers.

Especially in the context of the highly complex and non-academic programming languages C and C++, which are the state of the art in system software and embedded software development, the implementation of optimizations based on static whole program analysis is extremely difficult. Although many of the aforementioned algorithms and techniques were proposed at least 10 years ago, today's commercial compilers have not incorporated these innovations yet. The main reasons are their complexity and existing de-facto standards like object file formats and ABIs.[9] Therefore, we pragmatically advocate for using OO and AO mechanisms in combination in order to get the best out of both worlds. While OO is well suited for applications that actually need dynamism, static advice weaving is an interesting alternative that avoids the OO overhead from the beginning if dynamism is not required.

5.3   Applicability to System Software

Compared to application software product lines, system software product lines have to fulfill even higher demands on configurability and resource-efficiency. With respect to efficiency, the aforementioned abstractions are particularly useful in system software:

---

[9] ABI = <u>A</u>pplication <u>B</u>inary <u>I</u>nterface.

efficient callbacks by advice code inlining. System software usually has to provide means to bind user-defined code to (asynchronous) system events, such as hardware interrupts in an operating system. To prevent uses-dependencies [41] of the system code to driver/application code, a *callback* or *upcall* mechanism via function-pointers has to be provided. This leads to extra overhead and higher latencies at run time. As aspects reverse such dependencies [14], they provide means to break user-dependencies without this extra overhead.

handler chains by multiple advice. The implementation of handler chains can be implemented by the same mechanism as single callbacks. Due to advice code inlining, this furthermore induces no additional overhead.

configurable initialization order by aspect precedences. In system software it is often necessary to initialize components and subsystems in some well-defined (partial) order during system startup. The order is determined by component-interdependencies; however, not necessarily fixed in the case of configurable product lines. It may well depend on the actual configuration. By aspect precedences (given as *order advice* in AspectC++), aspects provide perfect means to solve this problem at compile time.

5.4 Potential Disadvantages of the Approach

The advantages regarding resource efficiency are mainly caused by binding and inlining component code at compile time that is bound at run time in OOP. This may, however, also lead to limitations of the approach:

compile time fixation limits the approach to static configuration of product lines. In domains, where the set of selected features can arbitrarily change at run time, additional support has to be provided, e.g., by a dynamic aspect weaver. As most dynamic weaving approaches induce inevitable overhead [27], it is questionable if in such cases AOP still outweighs the overhead of OOP. However, in most cases only some features can be expected to change at run time. A tailorable low-cost weaving infrastructure, which provides means to configure *per-feature* whether static or dynamic weaving should be used [45], may be a promising solution for such cases, as this enables developers to optimize the trade-off between dynamism and resource overhead at configuration time.

code bloating due to advice code inlining. This occurs if functionality given by advice is invoked from (and therefore woven to) many different joinpoints. While this situation has to be considered for highly crosscutting concerns such as tracing, component composition usually takes place at few well-known joinpoints only. In any case, however, the programmer is able to prevent this effect by simply moving the advice body into a non-inlined function that is invoked from the advice code.

5.5 Applicability to Other AOP Approaches

Our results show, that in the case of C++ and AspectC++, using AOP instead of OOP may lead to significant saving effects in terms of memory and hardware costs. This

gives rise to the question, whether similar results can be reached with other languages and/or AOP approaches as well. This is clearly a huge field for further research and cannot be answered in this paper. We think, however, that the following points are of particular importance:

weaver run time support may lead to some unavoidable base overhead that limits the reachable quantity of memory savings. This is clearly an issue for the domain of embedded systems. Experiments conducted with AspectJ on Java2 Micro Edition, for instance, have shown that the AspectJ run time induces an extra 35 KB overhead [49]. AOP approaches that do not require run-time library support are probably better suited for the goal of resource-thriftiness. Besides AspectC++, XWeaver [44] and AspectC [12], for instance, may be promising candidates.

language capabilities of the host language have a high influence on the applicability of the approach. As the resource savings are mainly caused by using compile time instead of run-time concepts, the host language has to support such concepts. On the one hand, it is probably difficult to get "rid of the OOP overhead" in languages that have been particularly designed as object-oriented languages, such as Java or C#. On the other hand, languages such as C may not be sufficiently expressive with respect to (compile time) genericity to take full advantage of the approach.

5.6   Design Issues

The actual design and implementation of the AO and OO versions have a significant influence on the resulting memory and performance numbers. Both versions have achieved the *granularity*, *pluggability*, and *extensibility* goals. One might, however, question if especially the design of the OO version is optimal with respect to the *economy* goal as well. Theoretically, it might be possible to find a better OO implementation that leads to lower memory demands.

As a matter of fact, it is impossible to proof that the design and implementation of some software is optimal with respect to run-time resource requirements. The authors have, however, profound experience in developing resource-minimal object-oriented product lines [10, 11, 34]. There is, furthermore, quite some evidence that any OO version would lead to higher memory requirements than the AO version: as pointed out in the previous sections, run-time mechanisms have to be used with OO to reach the required flexibility which, as described in Sects. 3.5 and 4, induce overhead. Given that the OO version takes up to 138% more RAM and up to 91% more flash memory than the C and AO versions; it is hard to imagine that these differences could be levelled by a better OO implementation.

For the AO version the resource requirements are only minimally higher than for the C version. Hence, it can already be considered as "nearly optimal". In a recent study on the eCos operating system product line [1], we could show that the C-mechanism to reach configurability (conditional compilation), can be replaced by AOP concepts without any extra run time and memory costs. In this study, we refactored several hundred `#ifdef`-blocks caused by configurable features into aspects [33].

6   Related Work

Several papers have been published that analyze and compare different implementation techniques of product line variabilities [23, 36]. All authors come to the conclusion that AOP is a very promising approach, although still not widely used in the industrial practice [38]. In the middleware area several case studies showed that AOP allows software developers to scale features and footprint by supporting the static configuration of fine-grained configurable features [13, 50, 51]. All of these studies were conducted with AspectJ. In one case the authors clearly expressed that AspectC++ would be more appropriate for the embedded systems domain [28].

Other proposed implementation techniques were Generative Programming [16], the application of frame processors such as XVCL [29] or Angie [19], and Feature-Oriented Programming (FOP) [6]. All three techniques have in common with AOP that in a layered system design the higher-level layers can use special means to refine the behavior and/or structure of the lower-level layers. This *dependency alignment* [14] is a key to successful product line design. In comparison to frame processing and C++ template-based generative programming [17], the main advantage of the AOP approach is its *obliviousness*, i.e., refined layers do not have to be prepared for intervention by aspects. In comparison with refinements in FOP, the advantage of AOP is *quantification*, i.e., a single aspect can affect various different joinpoints. Of course, FOP is much more than just the refinement mechanism. Therefore, the novel FeatureC++ language [5] is very promising, as it combines AOP and FOP.

When using OOP, the discussed refinement can only be achieved with virtual functions at the cost of dynamic binding. However, in many cases dynamic binding offers more flexibility than it is actually needed. Much work has been conducted to reduce the costs of virtual functions [3, 8, 20], but these optimizations are still not state of the art in current C++ compilers.

Another alternative approach is to combine OOP with partial program evaluation tools like Tempo [15] and C-Mix [4], especially as partial evaluation has succesfully been applied in the context of system software [37]. However, the authors are not aware of any tools or concept papers that cover the combination of these techniques with OOP. Furthermore, the case studies on partial evaluation of system software focus on performance improvements. Some of these techniques increase the code size by keeping specialized and original code in memory at the same time or by applying dynamic compilation [42].

In the systems software area configurable software is often component based. Popular examples are the OSKit [22, 39] and TinyOS [7]. In both cases the requirements on performance and footprint can only be achieved with additional tool support. For example, Knit can be regarded as a weaver for OSKit components [43]. TinyOS is written in the NesC language [25]. Here, the compiler performs a whole program transformation in order to optimize the system. In both cases the overhead, which is the result of the component model, is later removed by an additional tool. With the AOP approach an overhead is avoided from the beginning.

An interesting alternative to layered designs is Subject-Oriented Programming (SOP) [26, 40]. This technique is one of the roots of AOP. It supports system composition from

slices instead of layers. In our embedded weather station case study, the sensors are in fact slices. By using the introduction mechanism of AspectC++, SOP was simulated.

7   Summary and Conclusions

The development of fine-grained and resource-efficient system software product lines requires means for separation of concerns that do not lead to an extra overhead in terms of memory and performance. Especially in the important domain of small embedded systems, where even a few additional bytes of code or data may lead to higher hardware costs, the inherent overhead of OOP is not acceptable for most developers. For this domain we propose AOP as an alternative.

In this paper, we compared an object-oriented with an aspect-oriented implementation of an configurable embedded software product line. We could show that the aspect-oriented implementation induces significantly lower memory requirements than the object-oriented implementation, while providing similar or even better qualities with respect to separation of concerns. Although our results were produced with C++ and AspectC++, similar effects should be achievable with any AOP approach that does not require additional run-time support and performs inlining of advice code. Advice code inlining also leads to an excellent performance. Our measurements have shown that the performance of the AO version was identical with the performance of our C-based reference implementation.

Overall, our results show that aspects provide very good separation of concerns in conjuction with high resource efficiency. This makes AOP well suited for the development of embedded system software product lines—where aspects beat objects.

Regarding future work, we will continue to evaluate costs and benefits of applying AOP to the operating systems domain. We are furthermore working on a benchmark suite that provides detailed results regarding the run time and memory costs of AspectC++ features.

Acknowledgments

References

[1] eCos homepage. http://ecos.sourceware.org/
[2] OSEK/VDX standard. http://www.osek-vdx.org/
[3] G. Aigner and U. Hölzle. Eliminating virtual function calls in C++ programs. Technical Report TRCS95-22, Computer Science Department, University of California, Santa Barbara, December 1995

[4] L.O. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD Thesis, DIKU, University of Copenhagen, May 1994. (DIKU report 94/19)

[5] S. Apel, T. Leich, M. Rosenmller, and G. Saake. FeatureC++: on the symbiosis of feature-oriented and aspect-oriented programming. In: *Proceedings of the 4th International Conference on Generative Programming and Component Engineering (GPCE '05)*, Tallinn, Estonia, 2005

[6] D. Batory. Feature-oriented programming and the AHEAD tool suite. In: *Proceedings of the 26th International Conference on Software Engineering (ICSE '04)*, IEEE Computer Society, pp. 702–703, 2004

[7] U.C. Berkeley. TinyOS homepage. http://www.tinyos.net/

[8] D. Bernstein, Y. Fedorov, S. Porat, J. Rodrigue, and E. Yahav. Compiler optimization of C++ virtual function calls. In: *Proceedings of the 2nd USENIX Conference on Object-Oriented Technologies and Systems (COOTS '96)*, Toronto, Canada, 1996

[9] D. Beuche. Variant management with pure::variants. Technical report, pure-systems GmbH, 2003. http://www.pure-systems.com/

[10] D. Beuche, A. Guerrouat, H. Papajewski, W. Schröder-Preikschat, O. Spinczyk, and U. Spinczyk. The PURE family of object-oriented operating systems for deeply embedded systems. In: *Proceedings of the 2nd IEEE International Symposium on Object-Oriented Real-Time Distributed Computing (ISORC '99)*, St Malo, France, pp. 45–53, 1999

[11] D. Beuche, W. Schröder-Preikschat, O. Spinczyk, and U. Spinczyk. Streamlining object-oriented software for deeply embedded applications. In: *Proceedings of the 33rd International Conference on Technology of Object-Oriented Languages and Systems (TOOLS '00)*, Mont Saint-Michel, France, pp. 33–44, 2000

[12] Y. Coady, G. Kiczales, M. Feeley, and G. Smolyn. Using AspectC to improve the modularity of path-specific customization in operating system code. In: *Proceedings of the 3rd Joint European Software Engineering Conference and ACM Symposium on the Foundations of Software Engineering (ESEC/FSE '01)*, 2001

[13] A. Colyer, A. Clement, R. Bodkin, and J. Hugunin. Using AspectJ for component integration in middleware. In: *Proceedings of the 18th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '03)*, ACM, New York, pp. 339–344, 2003

[14] A. Colyer, A. Rashid, and G. Blair. On the separation of concerns in program families. Technical Report COMP-001-2004, Lancaster University, 2004

[15] C. Consel, L. Hornof, R. Marlet, G. Muller, S. Thibault, and E.-N. Volanschi. Tempo: specializing systems applications and beyond. *ACM Computing Surveys*, 30(3es), 1998

[16] K. Czarnecki and U.W. Eisenecker. *Generative Programming. Methods, Tools and Applications.* Addison-Wesley, Reading, 2000

[17] K. Czarnecki and U.W. Eisenecker. Synthesizing Objects. In: R. Guerraoui (ed.), *Proceedings of the 13th European Conference on Object-Oriented Programming (ECOOP '99)*, *LNCS vol. 1628*, Lisbon, Portugal, Springer, Berlin Heidelberg New York, pp. 18–42, 1999

[18] J. Dean, C. Chambers, and D. Grove. Selective specialization for object-oriented languages. In: *Proceedings of PLDI '95*, La Jolla, CA, 1995

[19] Delta Software Technology GmbH. Angie – an introduction, 2005

[20] K. Driesen and U. Hölzle. The direct cost of virtual function calls in C++. In: *Proceedings of the 11th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '96)*, 1996

[21] B. Dufour, C. Goard, L. Hendren, C. Verbrugge, O. de Moor, and G. Sittampalam. Measuring the dynamic behaviour of AspectJ programs. In: *Proceedings of the 19th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '04)*, ACM, New York, pp. 150–169, 2004.

[22] B. Ford, G. Back, G. Benson, J. Lepreau, A. Lin, and O. Shivers. The flux OSKit: a subst-rate for Kernel and language research. In: *Proceedings of the 17th ACM Symposium on Operating Systems Principles (SOSP '97), ACM Operating Systems Review*, ACM, New York, pp. 38–51, 1997

[23] C. Gacek and M. Anastasopoules. Implementing product line variabilities. In: *Proceedings of 2001 Symposium on Software Reusability: Putting Software Reuse in Context*, ACM, New York, pp. 109–117, 2001

[24] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley, Reading, 1995

[25] D. Gay, P. Levis, R. von Behren, M. Welsh, E. Brewer, and D. Culler. The nesC language: a holistic approach to networked embedded systems. In: *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '03)*, ACM, San Diego, New York, pp. 1–11, 2003

[26] W. Harrison and H. Ossher. Subject-oriented programming—a critique of pure objects. In: *Proceedings of the 8th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '93)*, pp. 411–428, 1993

[27] M. Haupt and M. Mezini. Micro-measurements for dynamic aspect-oriented systems. In: *NetObjectDays (NODe '04)*, Erfurt, Germany, *LNCS vol. 3263*, Springer, Berlin Heidelberg New York, pp. 81–96, 2004

[28] F. Hunleth and R. Cytron. Footprint and feature management using aspect-oriented pro-gramming techniques. In: *Proceedings of the 2002 Joint Conference on Languages, Com-pilers and Tools for Embedded Systems & Soft. and Compilers for Embedded Systems (LCTES/SCOPES '02)*, Berlin, Germany, ACM, New York, pp. 38–45, 2002

[29] S. Jarzabek and H. Zhang. XML-based method and tool for handling variant requirements in domain model. In: *Proceedings of the 5th IEEE International Symposium on Require-ments Engineering (RE '01)*, Toronto, Canada, IEEE Computer Society Press, pp. 116–123, 2001

[30] J.M. Kahn, R.H. Katz, and K.S.J. Pister. Next century challenges: mobile networking for "smart dust". In: *International Conference on Mobile Computing and Networking (MOBI-COM '99)*, pp. 271–278, 1999

[31] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W.G. Griswold. An overview of AspectJ. In: J. Lindskov Knudsen (ed.), *Proceedings of the 15th European Conference on Object-Oriented Programming (ECOOP '01), LNCS vol. 2072*, Springer, Berlin Heidelberg New York, pp. 327–353, 2001

[32] D. Lohmann, G. Blaschke, and O. Spinczyk. Generic advice: on the combination of AOP with generative programming in AspectC++. In: G. Karsai and E. Visser (ed.), *Proceed-ings of the 3rd International Conference on Generative Programming and Component En-gineering (GPCE '04)*, *LNCS vol. 3286*, Springer, Berlin Heidelberg New York, pp. 55–74, 2004

[33] D. Lohmann, F. Scheler, R. Tartler, O. Spinczyk, and W. Schröder-Preikschat. A quantita-tive analysis of aspects in the eCos kernel. In: *Proceedings of the EuroSys 2006 Conference (EuroSys '06)*, ACM, New York, pp. 191–204, 2006

[34] D. Lohmann, W. Schröder-Preikschat, and O. Spinczyk. On the design and development of a customizable embedded operating system. In: *Proceedings of the SRDS Workshop on Dependable Embedded Systems (SRDS-DES '04)*, IEEE Computer Society, pp. 1–6, 2004

[35] D. Lohmann and O. Spinczyk. On typesafe aspect implementations in C++. In: F. Geschwind, U. Assmann, and O. Nierstrasz (eds.), *Proceedings of Software Composi-tion 2005 (SC '05)*, Edinburgh, UK, *LNCS vol. 3628*, Springer, Berlin Heidelberg New York, pp. 135–149, 2005

[36] M. Mezini and K. Ostermann. Variability management with feature-oriented programming and aspects. In: *Proceedings of ACM SIGSOFT '04 / FSE-12*, 2004

[37] G. Muller, E.N. Volanschi, and R. Marlet. Scaling up partial evaluation for optimizing the sun commercial rpc protocol. In: *ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, Amsterdam, The Netherlands, ACM, New York, pp. 116–125, 1997

[38] D. Muthig and T. Patzke. Generic implementation of product line components. In: *NetObjectDays (NODe '02)*, Erfurt, Germany, *LNCS vol. 2591*, Springer, Berlin Heidelberg New York, pp. 313–329, 2003

[39] University of Utah. OSKit homepage. http://www.cs.utah.edu/flux/oskit

[40] H. Ossher and P. Tarr. Using multidimensional separation of concerns to (re)shape evolving software. *Communications of the ACM*, pp. 43–50, 2001

[41] D.L. Parnas. Some hypothesis about the uses hierarchy for operating systems. Technical report, TH Darmstadt, Fachbereich Informatik, 1976

[42] C. Pu, H. Massalin, and J. Ioannidis. The Synthesis kernel. *Computing Systems*, 1(1): 11–32, 1988

[43] A. Reid, M. Flatt, L. Stoller, J. Lepreau, and E. Eide. Knit: component composition for systems software. In: *Proceedings of the 4th Symposium on Operating Systems Design and Implementation (OSDI '00)*, San Diego, CA, USA, pp. 347–360, 2000 Usenix Association

[44] O. Rohlik, A. Pasetti, V. Cechticky, and I. Birrer. Implementing adaptability in embedded software through aspect oriented programming. In: *Proceedings of Mechatronics & Robotics (MechRob '04)*, Aachen, Germany, IEEE Computer Society Press, 2004

[45] W. Schröder-Preikschat, D. Lohmann, W. Gilani, F. Scheler, and O. Spinczyk. Static and dynamic weaving in system software with AspectC++. In: Y. Coady, J. Gray, and R. Klefstad (eds.), *Proceedings of the 39th Hawaii International Conference on System Sciences (HICSS '06) - Mini-Track on Adaptive and Evolvable Software Systems*. IEEE Computer Society, 2006

[46] O. Spinczyk, D. Lohmann, and M. Urban. Advances in AOP with AspectC++. In: H. Fujita and M. Mejri (eds.), *New Trends in Software Methodologies, Tools and Techniques (SoMeT '05), LNCS vol. 129*, Frontiers in Artificial Intelligence and Applications, Tokyo, Japan, IOS, pp. 33–53, 2005

[47] D. Tennenhouse. Proactive computing. *Communications of the ACM*, pp. 43–45, 2000

[48] M. Weiser. The computer for the 21st centrury. *Scientific American*, 265(3):94–104, 1991

[49] T. Young and G. Murphy. Using AspectJ to build a product line for mobile devices. *AOSD '05 Demo*, Chicago, Illinois, March 2005

[50] C. Zhang, D. Gao, and H.-A. Jacobsen. Generic middleware substrate through modelware. In: *Proceedings of the ACM/IFIP/USENIX 6th International Middleware Conference (Middleware '05)*, Grenoble, France, ACM, New York, 2005

[51] C. Zhang and H.-A. Jacobsen. Refactoring middleware with aspects. *IEEE Transactions on Parallel and Distributed Systems*, 14(11), 2003