

Machine-checked object layout for C++ multiple inheritance with empty-base optimization

Tahina Ramananandro

<http://gallium.inria.fr/~tramanan/cpp/object-layout>

August 6, 2010

Notations and conventions

We adopt the usual Coq list notations:

- `nil` is the empty list
- `a :: q` is the list starting with an element `a` and continuing with the tail list `q`
- `+` is the *append* (list concatenation) operator: for any list `l`, we have $\text{nil} + l \stackrel{\text{def.}}{=} l$ and $(a :: q) + l \stackrel{\text{def.}}{=} a :: (q + l)$.

We also adopt the following additional notations:

- `first` is a function defined on non-empty lists, such that $\text{first}(a :: l') \stackrel{\text{def.}}{=} a$ for all `a, l'`.
- `last` is a function defined on non-empty lists, computing their last elements: $\text{last}(a :: \text{nil}) \stackrel{\text{def.}}{=} a$ and $\text{last}(a :: b :: l') \stackrel{\text{def.}}{=} \text{last}(b :: l')$ for all `a, b, l'`.
- `length(l)` is the length of a list `l`: $\text{length}(\text{nil}) \stackrel{\text{def.}}{=} 0$ and $\text{length}(a :: l') \stackrel{\text{def.}}{=} 1 + \text{length}(l')$
- For any sets \mathcal{A} and $\mathcal{S} \subseteq \mathcal{A} \times \mathbb{Z}$, and any $\omega \in \mathbb{Z}$:

$$\omega + \mathcal{S} \stackrel{\text{def.}}{=} \{(A, \omega + o) \mid (A, o) \in \mathcal{S}\}$$

- For any two sets \mathcal{S}, \mathcal{T} : $\mathcal{S} \perp \mathcal{T}$ stands for $\mathcal{S} \cap \mathcal{T} = \emptyset$
- for any set \mathcal{S} , we pose

$$\text{option } \mathcal{S} \stackrel{\text{def.}}{=} \{\emptyset\} \cup \{\{S\} \mid S \in \mathcal{S}\}$$

which is the set of all subsets of \mathcal{S} with cardinality 0 or 1.

- for any two integers u, v : $(u \mid v)$ if and only if u *evenly divides* v , i.e. v is a multiple of u :

$$(u \mid v) \Leftrightarrow \exists k : v = k \cdot u$$

Our examples are given either by GNU GCC 4.3 on an Intel 32-bit machine architecture (where `sizeof(int) = 4`) with option `-malign-double` enabled, or by the prescriptions of Itanium C++ ABI [1].

Contents

1	Goal	2
2	The semantics of C++ multiple inheritance	2
2.1	Class hierarchy (<code>Cplusplus.v</code>)	2
2.2	Inheritance paths, or Base class subobjects (<code>Cplusplus.v</code>)	4
2.2.1	Non-virtual inheritance	4
2.2.2	Virtual inheritance	5
2.2.3	A practical example	6
2.3	Structure array fields	8
2.4	Well-founded hierarchies (<code>Cplusplus.v</code>)	11
3	Formalization of object layout	12
3.1	Methodological overview	12
3.1.1	The dynamic type data of a subobject (or “pointer to virtual table”)	12
3.1.2	Layout scheme	13
3.1.3	The need for an empty base optimization	13
3.1.4	Sizes, data sizes, and their non-virtual counterparts	14
3.1.5	Alignments and non-virtual alignments	15
3.2	The expected output of a layout algorithm	16
3.3	Soundness	19
3.3.1	Total size	19
3.3.2	Alignment	21
3.3.3	Data size	25
3.3.4	Non-overlapping of data	28
3.3.5	Dynamic type data	33
3.3.6	Identity of subobjects	39
4	Implementation of realistic layout algorithms (<code>CPP.v</code>, <code>CommonVendorABI.v</code>, <code>CCPP.v</code>)	44
4.1	An algorithm based on the Itanium C++ ABI	45
4.2	An optimized algorithm	49

1 Goal

Our work aims at:

- formalizing the object-oriented part of C++ with multiple inheritance and structure arrays
- finding realistic implementation criteria for C++ object layout
- proving that any implementation matching our criteria is correct with respect to our formalization

Our work is machine-checked with the Coq proof assistant.

More context information about the motivation of our work can be found in our POPL 2011 draft [6].

2 The semantics of C++ multiple inheritance

We started from [9] who formalized C++ multiple inheritance in Isabelle based on [7]. We extend their work by featuring structures and structure arrays.

2.1 Class hierarchy (`Cplusplusconcepts.v`)

Notation 1. Let \mathcal{A} be the set of atomic types (*int, float, short, ...*).

We make no further assumption on atomic types or values: our formalization is a Coq functorial module taking them as parameters.

Definition 1. A class hierarchy is a tuple $(\mathcal{C}, \mathcal{FI}, \mathcal{ScF}, \mathcal{StF}, \mathcal{DNV}, \mathcal{DV})$, where:

- \mathcal{C} is the set of classes. We assume $\mathcal{C} \cap \mathcal{A} = \emptyset$.
- \mathcal{FI} is the set of field identifiers
- $\mathcal{ScF} : \mathcal{C} \rightarrow \mathfrak{P}(\mathcal{FI} \times (\mathcal{A} \cup \mathcal{C}))$ gives, for each class, the set of its scalar fields (a scalar field is either an atom, or a pointer to C for some class C).
- $\mathcal{StF} : \mathcal{C} \rightarrow \mathfrak{P}(\mathit{StructField})$ (where $\mathit{StructField} \stackrel{\text{def.}}{=} \mathcal{FI} \times \mathcal{C} \times \mathbb{N}^*$)¹ gives, for each class, the set of its structure array fields. Each structure array field (f, C, n) is considered to be an array of n structures of type C . In fact, structure fields are structure array fields where $n = 1$.

¹The Standard forbids arrays with zero cells.

- $\mathcal{DNV} : \mathcal{C} \rightarrow \mathfrak{P}(\mathcal{C})$ gives, for each class, the set² of its non-virtual direct bases.
- $\mathcal{DV} : \mathcal{C} \rightarrow \mathfrak{P}(\mathcal{C})$ gives, for each class, the set of its virtual direct bases.

For instance, the following code:

```

struct A          { int i; };
struct B: virtual A { C * c; };
struct C: A, B    { float j; A a[2]; B b; };

```

would give the following hierarchy:

$$\begin{array}{l}
 \mathcal{C} = \{A, B, C\} \\
 \hline
 \mathcal{FI} = \{i, j, a, b, c\} \\
 \hline
 \begin{array}{l}
 A \mapsto \{(i, \text{int})\} \\
 \text{Sc}\mathcal{F} : B \mapsto \{(c, C)\} \\
 C \mapsto \{(j, \text{float})\}
 \end{array} \\
 \hline
 \begin{array}{l}
 A \mapsto \emptyset \\
 \text{St}\mathcal{F} : B \mapsto \emptyset \\
 C \mapsto \{(a, A, 2), (b, B, 1)\}
 \end{array} \\
 \hline
 \begin{array}{l}
 A \mapsto \emptyset \\
 \mathcal{DNV} : B \mapsto \emptyset \\
 C \mapsto \{A, B\}
 \end{array} \\
 \hline
 \begin{array}{l}
 A \mapsto \emptyset \\
 \mathcal{DV} : B \mapsto \{A\} \\
 C \mapsto \emptyset
 \end{array}
 \end{array}$$

If a Coq formalization had to exactly mimic these definitions, then it would need dependent types, as the domain of, say, \mathcal{DNV} depends on \mathcal{C} , both being part of the hierarchy data. But we chose to avoid dependent types to represent those functions and families, by using finite mappings and externally imposing *well-formedness* constraints (`CplusWf.v`).

In practice, `Compcert` provides a library for finite mappings (`Maps.v`), to represent functions $f : \mathbb{N}^* \rightarrow \text{option } \mathcal{S}$ such that the “domain” $\{n \in \mathbb{N}^* \mid f(n) \neq \emptyset\}$ is finite. For any set \mathcal{S} , the set of such finite mappings is denoted `PTree.t S` in respect to the tree implementation.

In our case, we do not explicitly define the set \mathcal{C} of the classes of a hierarchy. Instead, we define (`Cplusconcepts.v`) a hierarchy to be a mapping $h \in \text{PTree.t Class.t}$ where `Class.t` is a record type having fields such as $\mathcal{DNV}, \mathcal{DV}, \dots$. Then, class identifiers are in \mathbb{N}^* and, morally, \mathcal{C} is the “domain” of h , that is the set of class identifiers C such that $h(C) \neq \emptyset$.

But we also avoid dependent types for the value types of the fields of `Class.t`. For instance, we have, for any class C such that $h(C) = \{c\}$ for some $c \in \text{Class.t}$, $c.\mathcal{DNV} \in \text{list } \mathbb{N}^*$ (and not `list C`, which would require dependent

² For our purpose, direct base ordering is not relevant.

types, as C depends on h). To enforce the domains, we externally require *completeness* conditions (`CplusplusWf.v`) such as:

$$\forall C, c : h(C) = \{c\} \Rightarrow \forall b \in c.\mathcal{DNV}, h(b) \neq \emptyset$$

which states that every direct non-virtual base of a class is a class belonging to the hierarchy. (More precisely, every identifier of the set of direct non-virtual base identifiers is a valid identifier of an existing class.)

2.2 Inheritance paths, or Base class subobjects (`Cplusplusconcepts.v`)

2.2.1 Non-virtual inheritance

Consider the following code:

```

struct A           { int a; };
struct B1 : A       {};
struct B2 : A       {};
struct C : A        {};
struct D : B1, B2, C {};

```

Then, an instance of D will have *three* different “copies” of A : one reachable through the direct non-virtual base $B1$, one through $B2$, and one from C . This is called *non-virtual inheritance*, or *repeated inheritance*. Each “copy”, called a *subobject* of D of static type A , has its own value for the field a . So, it is necessary to distinguish those subobjects. Following Wasserrab et al., those subobjects can be distinguished by the *paths* through which they are reached.

More precisely:

Definition 2. A list l is a non-virtual path from C to A if, and only if:

- either $C = A$ and $l = A :: \text{nil}$. This path is called the trivial path.
- or there exists a non-virtual direct base B of C and a non-virtual path l' from B to A , such that $l = C :: l'$.

A is a non-virtual base of C if, and only if, A is reachable through a non-trivial non-virtual path from C .

Then, in the above example, A is a non-virtual base of C through two *different* paths: $C :: B1 :: A :: \text{nil}$ and $C :: B2 :: A :: \text{nil}$. But $B2$ is not a non-virtual base of $B1$.

Then, if we have

```

C      c;
B1 * b1 = (B1 *) &c;
B2 * b2 = (B2 *) &c;
A  * a1 = (A  *) b1;
A  * a2 = (A  *) b2;

```

the C++ Standard [4] dictates $a_1 \neq a_2$: even though they are pointers to subobjects of the same C object, those subobjects are accessible through *different non-virtual paths*, so those pointers must be different.

This is one of the properties our work aims at proving.

In particular:

Lemma 1. *Any non-virtual path from any class C to any class B begins with C and ends with B .*

which leads to the:

Lemma 2. *If $C :: l$ is a non-virtual path from C to B and $B :: l'$ is a non-virtual path from B to A , then $C :: l+l'$ is a non-virtual path from C to A denoted $(C :: l)@_{Repeated}(B :: l')$.*

2.2.2 Virtual inheritance

C++ features *virtual inheritance*, to allow some subobjects accessible through apparently different paths to be equal. Rossie and Friedman [7] (formalized in Isabelle by Wasserrab, Nipkow et al. [9]) give a notion of path allowing such subobjects to be represented with actually *the same* path.

Definition 3. *A class A is a virtual base of C if, and only if:*

- *either A is a direct virtual base of C*
- *or there is a direct (virtual or non-virtual) base B of C such that A is a virtual base of B .*

Definition 4. *A path, inheritance path, or base class subobject, from a class C to a class A is a couple (h,l) such that:*

- *either $h = Repeated$ and l is a non-virtual path from C to A*
- *or $h = Shared$ and there exists a virtual base B of C such that l is a non-virtual path from B to A .*

The set of all paths shall be denoted $Path$.

For instance, if we have:

```

struct A           {};
struct B1: virtual A  {};
struct B2: virtual A  {};
struct C : A         {};
struct D : A, B1, B2, C {};

```

Then, the path from D to A through B1 denotes the same subobject as the path through B2, because A is a *virtual* base of both B1 and B2: the A subobject is *shared* by both classes. As B1 and B2 are bases of D, then A is a virtual base of D and this subobject is denoted by the path (Shared, A :: nil). However, A is also accessible from D through two non-virtual paths: $D :: A :: \text{nil}$ (because A is a non-virtual direct base of D), and $D :: C :: A :: \text{nil}$. Those paths denote three different subobjects.

Lemma 3. *Let C_0, C_1, C_2 be three classes. If (h_i, l_i) is a path from C_{i-1} to C_i for each $i \in 1, 2$, then:*

- *if $h_2 = \text{Repeated}$, then $(h_1, l_1 @_{\text{Repeated}} l_2)$ is a path from C_0 to C_2*
- *if $h_2 = \text{Shared}$, then (h_2, l_2) is a path from C_0 to C_2*

This path is denoted $(h_1, l_1) @ (h_2, l_2)$.

2.2.3 A practical example

Non-virtual and virtual multiple inheritance are not uncommon in everyday life. Consider for instance three electronic devices: a **radio**, an **alarm**, and a **radio-alarm**.

- A *radio* needs electronic current provided by a **plug**. It has a **switch** to turn it on or off. The radio station currently running can be controlled via its **frequency**.
- An *alarm* needs electronic current provided by a **plug**. It has a **switch** to turn it on or off. The **wake-up time** tells when the alarm should ring up.
- A *radio-alarm* combines both a **radio** and **alarm**. Either the **radio** and the **alarm** can be turned on or off independently of the other, through the corresponding **switch**. By contrast, the radio-alarm also needs electric current provided by a **plug** common to both radio and alarm components.

A radio-alarm device could be modelled by a C++ program using multiple inheritance, and especially both non-virtual and virtual inheritance.

```

struct Plug {
    bool plugged;
}

struct Device: virtual Plug {
    bool switch;
}

struct Radio: Device {

```

```

    double frequency ;
}

typedef double GNUEpoch;

struct Alarm : Device {
    GNUEpoch wakeUpTime;
    virtual void ring ();
}

struct RadioAlarm : Radio , Alarm {
    void ring () {
        if (Radio::switch) {
            Alarm::ring ();
        } else {
            Radio::switch = true;
        }
    }
}

```

Here, a RadioAlarm will have one frequency for its Radio, one wakeUpTime for its Alarm, two **switches** (one for its Radio, one for its Alarm), but only one plugged status. It is impossible to unplug the Radio without unplugging the Alarm: they share the same plug.

The Alarm has its own method to ring up, usually by buzzing. This method is called at timeout, through the following code:

```

GNUEpoch getTime ();

void wait (Alarm * alarm) {
    if (getTime () > alarm->wakeUpTime) return ;
    while (getTime () < alarm->wakeUpTime) {};
    alarm->ring ();
}

```

But a RadioAlarm may want to turn the radio on instead of buzzing, unless the radio is already turned on. That is why the RadioAlarm class *redefines*, or *overrides*, the method ring, so that the following code:

```

int main (int , String []) {
    RadioAlarm a;
    wait (static_cast<Alarm*>(&a));
    return 0;
}

```

will actually call the ring method defined in class RadioAlarm. To enable this, ring must be declared **virtual** in Alarm.

2.3 Structure array fields

Our work extends Wasserrab, Nipkow et al. with *structure array fields*. That is, *whole objects* can be embedded in a given structure field or structure array field. Keeping this point of view would break the condition of two whole objects to be disjoint, so it is necessary to take structure fields into account by formalizing a notion of *generalized subobjects* including structure field access, not only inheritance.

Our work features arrays, but only considers structure arrays. Indeed, arrays of scalars can be simulated through arrays of structures having only one scalar field. For instance³:

```
int * t = new int[18];
t[4] = 42;
```

can be simulated with:

```
struct Sint { int value; };
Sint * t = new Sint() [18];
t[4].value = 42;
```

This point of view has no significant impact on performance⁴: accessing the field of a structure from a structure array yields only one memory access, as many as accessing a field from an array of scalars.

Any non-null object pointer b of type B^* evaluates to a pointer of the following (informal) form:

$$\begin{array}{rcl}
 C_1 * c_1 & = & ((B_1^*)\&c_0[i_1]) \rightarrow f_1 \\
 C_2 * c_2 & = & ((B_2^*)\&c_1[i_2]) \rightarrow f_2 \\
 & & \vdots \\
 C_k * c_k & = & ((B_k^*)\&c_{k-1}[i_k]) \rightarrow f_k \\
 B * b & = & (B^*)\&c_k[i]
 \end{array}$$

where:

- c_0 is a structure array of some class C_0 of size n_0
- $0 \leq k$
- for each $j \in [1 \dots k]$, $0 \leq i_j < n_{j-1}$
- for each $j \in [1 \dots k]$, there is an inheritance path P_j from C_{j-1} to some class B_j

³This simulation pattern works not only for atomic types, but also for pointer types. A generic pattern for scalar types would use a template. However, templates are out of the scope of our work.

⁴As long as Sint has only a default constructor, which may be ignored at execution. This is out of the scope of our work.

- for each $j \in [1 \dots k]$, there is a class C_j such that f_j is an structure array field of B_j of type $C_j[n_j]$ for some n_j
- $0 \leq i < n_k$
- there is an inheritance path P from C_k to B

We call the sequence $(i_1, P_1, f_1); (i_2, P_2, f_2); \dots; (i_k, P_k, f_k)$ an *array path* from $C_0[n_0]$ to $C_k[n_k]$. Let us denote it \mathcal{P} .

More formally:

Definition 5. Let $C, C' \in \mathcal{C}$, $n, n' \in \mathbb{N}$ and l be a list of elements of the set $\mathbb{N} \times \text{Path} \times \text{StructField}$. We say that l is an array path from $C[n]$ to $C'[n']$ if, and only if, one of these conditions holds:

- $C = C'$ and $n' \leq n$ and $l = \text{nil}$
- or all the following conditions hold:
 - there exists $i \in \mathbb{N}$ such that $i < n$
 - there is a class $A \in \mathcal{C}$ and a path $p \in \text{Path}$ from C to A
 - there is a structure array field $F = (f, D, s) \in \text{St}\mathcal{F}(A)$
 - there is a list l' such that l' is an array path from $D[s]$ to $C'[n']$
 - and $l = (i, p, F) :: l'$

Definition 6. A generalized subobject or relative pointer p of static type A relatively to an array of type $C[n]$ (or a relative pointer from $C[n]$ to A) is a triple (\mathcal{P}, i, P) where:

- \mathcal{P} is an array path from $C[n]$ to some $C'[n']$
- i is an index in the array of type $C'[n']$, such that $0 \leq i < n'$
- and P is an inheritance path from C' to A

Definition 7. A run-time pointer π of static type A is a couple (c, p) where:

- c is the initial structure array, either statically declared, or dynamically created, of some type $C[n]$
- p is a relative pointer from $C[n]$ to A

For instance, if we have the following code:

```

struct A0          {};
struct A : A0      {};
struct B          { A a[4]; };
struct C1: virtual B {};
struct C2: virtual B {};
struct D : C1, C2  {};
struct E          { D d[5]; };
E e[7];

```

then the expression:

$(A0 \ *) \ \&(e[2] \cdot d[0] \cdot a[3])$

yields the pointer:

(
 $e,$
 $(2, (\text{Repeated}, E :: \text{nil}), d) :: (0, (\text{Shared}, B :: \text{nil}), a) :: \text{nil},$
 $3,$
 $(\text{Repeated}, A :: A0 :: \text{nil})$
)

where $(2, (\text{Repeated}, E :: \text{nil}), d) :: (0, (\text{Shared}, B :: \text{nil}), a) :: \text{nil}$ is an array path from $E[7]$ to $A[4]$.

This apparently counterintuitive presentation allows us however to easily formalize object pointer arithmetics. Indeed, any *full object* (that is, any subobject that is not the result of a non-trivial inheritance cast) actually always belongs to some structure array (either initial or defined by a structure array field). So, if p is a pointer to such a full object, then $p[q]$ may be defined (not necessarily for any q however).

As p points to a full object, we have $B = C_k$, so the path P is trivial, and we have:

$$p = \&(c[i_1] \cdot f_1[i_2] \cdot f_2 \dots f_{k-1}[i_k] \cdot f_k[i]) = (c, \mathcal{P}, i, P)$$

so that, for any q such that $0 \leq i_{k+1} + q < n_k$:

$$p[q] = \&(c[i_1] \cdot f_1[i_2] \cdot f_2 \dots f_{k-1}[i_k] \cdot f_k[i + q]) = (c, \mathcal{P}, i + q, P)$$

That is, the only change is the index in the array of type $C_k[n_k]$. In particular, the array path and the inheritance path remain unchanged.

Our idea also allows to easily express the other two basic operations on pointers:

- if $p = (c, \mathcal{P}, i, P)$ is a pointer to some subobject of type B and if Q is an inheritance path from B to A such that derived-to-base casting of a pointer to B yields a pointer to A using Q , then $(A*)p = (c, \mathcal{P}, i, P@Q)$. Similarly for base-to-derived casting, if $p = (c, \mathcal{P}, i, P@Q)$, then base-to-derived casting would lead to (c, \mathcal{P}, i, P) .
- if $p = (c, \mathcal{P}, i, P)$ is a pointer to some subobject of type B and if f is a structure field of type A defined in class B , then $p \rightarrow f = (c, (\mathcal{P} \# (i, P, f) :: \text{nil}, 0, Q))$ where $Q = (\text{Repeated}, A :: \text{nil})$ is the trivial inheritance path from A to A . In other words, $p \rightarrow f$ points to the first structure of the array defined by f .

2.4 Well-founded hierarchies (CplusWf.v)

We expect some well-foundedness hypotheses to hold on the class hierarchy, so as to be able to *compute*, for any class C , some data $F(C)$ depending on $F(B)$ for all classes B being bases of C or types of the structure array fields of C .

Notation 2. For any class C , $\mathcal{D}(C)$ shall denote the set of its direct bases (virtual or non-virtual):

$$\mathcal{D}(C) \stackrel{\text{def.}}{=} \mathcal{DNV}(C) \cup \mathcal{DV}(C)$$

Hypothesis 1. Assume the hierarchy is well-founded, i.e. the set of classes can be ordered by some well-founded relation \prec compatible with inheritance, i.e. such that:

$$\forall B \in \mathcal{D}(C) : B \prec C$$

and compatible with structure array fields, i.e. such that:

$$\forall (f, B, n) \in \text{StF}(C) : B \prec C$$

In our Coq formalization, $\mathcal{C} \subseteq \mathbb{N}^*$, and \prec is the usual ordering relation on \mathbb{N}^* .

Notation 3. For each class C , the set of its virtual bases is denoted $\mathcal{V}(C)$. Thus, as the hierarchy is well-founded, we can compute $\mathcal{V}(C)$ by well-founded induction over \prec :

$$\mathcal{V}(C) \stackrel{\text{def.}}{=} \mathcal{DV}(C) \cup \bigcup_{B \in \mathcal{D}(C)} \mathcal{V}(B)$$

Lemma 4.

$$\forall C, \forall B \in \mathcal{V}(C) : B \prec C$$

In particular, $C \notin \mathcal{V}(C)$.

Lemma 5. If there exists a path from C to A , then $A \preceq C$.

Lemma 6. The only path from a class C to itself is the trivial path:

$$(\text{Repeated}, C :: \text{nil})$$

Lemma 7. For any classes C and A , if $(h, B :: l)$ is a path from C to A , then $B = C$ if and only if $h = \text{Repeated}$.

3 Formalization of object layout

In this section, the hierarchy is assumed to be well-founded.

Our goal is to compute, for any pointer p to a subobject of type A , a memory offset:

- such that if q is a pointer to a different subobject of type A , then p and q must have different memory offsets
- and, for any scalar field f defined in the class A , the offset in memory where the value of $p \rightarrow f$ is stored, such that any two different *scalar* fields never overlap. This allows showing the *good variable property* on concrete memory models such as CompCert.

Those two goals do not necessarily entail each other. Indeed, if A has no fields, then pointers to distinct subobjects of type A still have to be distinguished, which forbids structures from occupying zero-sized storage. To avoid waste of space, most compilers introduce *empty base optimizations* to allow empty structures of *different* types to overlap.

Our aim is to prove the correctness of a family of layout algorithms. So, we first give soundness conditions, which we prove to be correct with respect to the two

3.1 Methodological overview

Our work formalizes a layout principle common to most current C++ compilers.

3.1.1 The dynamic type data of a subobject (or “pointer to virtual table”)

When a virtual method is called from a subobject of some type A , most compilers use specific data such as a virtual table to call the method. This data is specific to the class of the whole object of the considered subobject. Indeed, the non-virtual structure of a subobject of type A is independent on how the subobject is reached: for any field of A (or of a non-virtual base of A), the offset to access this field will be always the same regardless of how the subobject is reached. A is the *static type* of the subobject. However, calling a method, or reaching a virtual base, depends on the way the subobject is reached: this additional information is the *dynamic type* of the subobject. At the semantics level, it is actually the inheritance path from the full object to the subobject.

At the implementation level, most C++ compilers introduce additional data at the beginning of some classes to store that dynamic type data. For instance, GNU GCC stores a pointer to the virtual table of the subobject, which is different depending on the dynamic type of the subobject. But

other compilers may use other additional data than virtual tables (for instance a key for dictionary search). Our formalization does not depend on the actual data stored, it only fixes the size.

In practice, such additional data are needed if the class has or inherits a virtual method, or has a virtual base. Such class is called *dynamic*.

However, our formalization of object layout does not rely on a precise definition of this notion, which is required only for concrete algorithms.

3.1.2 Layout scheme

We formalize the following layout for a full instance of some class C , common to many C++ compilers:

- The non-virtual part of C , consisting of:
 - The non-virtual part of a dynamic non-virtual direct base A of C , if such class exists
 - If C is dynamic but such A does not exist, then some space for the *dynamic type data* for C
 - Then the non-virtual parts of other non-virtual direct bases of C
 - Then the fields of C , so that they never overlap
- Then the non-virtual parts of all virtual bases of C , so that they never overlap

3.1.3 The need for an empty base optimization

Empty classes must have a nonzero size. Indeed, if for instance $\text{sizeof}(A) = 0$, then two objects from two distinct indices of an array of A could not be distinguished.

So assume $\text{sizeof}(A) = 1$. Then, it would be sound to have $\text{sizeof}(B_1) = 1$ and $\text{sizeof}(B_2) = 1$ for classes A, B_1, B_2 below:

```

struct A           { };
struct B1: A       { };
struct B2: A       { };
struct C : B1, B2 { int i; };
```

A naive compiler would lay out B_1 , B_2 and i in disjoint memory locations, yielding $\text{sizeof}(C) \geq 2 + \text{sizeof}(\text{int})$, that is, with alignment constraints, $\text{sizeof}(C) = 2 * \text{sizeof}(\text{int})$. However, it is consistent with the Standard to allow i to overlap empty bases, which would lead to $\text{sizeof}(C) = \text{sizeof}(\text{int})$, which is exactly the result given by GNU G++. This optimization is called *empty base optimization*: empty bases can overlap the data of non-empty bases or fields.

Though, offsets for B_1 and B_2 are still distinct, as each has a subobject of type A : if they were not, then we would again have $a_1 = a_2$ in the program below:

```

C    c ;
B1 * b1 = (B1 *) &c ;
B2 * b2 = (B2 *) &c ;
A  * a1 = (A  *) b1 ;
A  * a2 = (A  *) b2 ;

```

So, two empty bases may overlap only with care. We shall see later in more detail the precise constraints to follow when letting empty bases overlap other data. This is the main reason why C++ object layout cannot be mapped to C structure representations as Compcert [2] or cfront [8] do.

However, our work also shows that it is possible, without breaking the semantics of a C++ program, to generalize this empty base optimization to empty structure members.

This is the reason why our formalization gives no precise definition for empty classes. Instead, some constraints are enforced:

Hypothesis 2. *A layout algorithm may define its own notions of empty and dynamic classes, subject to the following constraints:*

- *An empty class must have no scalar fields*
- *An empty class must have no non-empty base*
- *An empty class must have no structure field of non-empty class type (regardless, however, of the number of array cells)*
- *A dynamic class must not be empty*

We shall see layout algorithms giving precise definitions for empty classes, sometimes more restrictive to lower compilation time, at the price of losing run-time optimizations on the programs being compiled.

3.1.4 Sizes, data sizes, and their non-virtual counterparts

Following the layout scheme above, we have to consider not only the size of a full object (which would be yielded by the `sizeof` operator), but also the size of the non-virtual parts of a class.

Moreover, dealing with empty classes that can overlap other components, considering their size or non-virtual size is not enough (as those sizes are necessary not null), we introduce the notion of *data size*. Roughly speaking, the *data size* of a class is an upper bound on the sum of the sizes of all scalar fields accessible from this class (through inheritance and/or structure array field paths). The *non-virtual data size* of a class is an upper bound on the sum of the sizes of all scalar fields accessible from this class *when the first inheritance step is non-virtual inheritance* (i.e. virtual inheritance is allowed only once a structure field within a non-virtual base of C is reached). Both sums also include the sizes of the dynamic type data for all relevant bases. The data size and non-virtual data size are not relevant for empty classes.

3.1.5 Alignments and non-virtual alignments

Consider the following hierarchy:

```
struct A {
    double d0;
    char c0;
};

struct B1: virtual A {
    int i1;
    char c1;
};

struct B2: virtual A {
    int i2;
};

struct C: B1, B2 {
};
```

GNU GCC 4.3 on Intel x86 32-bit platform with option `-malign-double` (to align **double** on 8-byte boundaries) gives the following layouts within C :

- B_1 is laid out at offset 0
- B_2 is laid out at offset 12
- A is laid out at offset 24

Even though B_2 inherits from A , it is laid out at offset 12 within C , which is not a multiple of 8 (the alignment of **double** $d0$). This shows that non-virtual bases B_1 and B_2 are aligned independently on the access to virtual base A . This leads us to distinguish two alignments for a class:

- the *alignment* of a class C , such that any instance of C be aligned at an offset multiple of this alignment
- the *non-virtual alignment* of a class C , to allow the alignment of C as a base of other classes, so that any subobject of type C be aligned at an offset multiple of this alignment

In our case, we can see that the non-virtual alignment of B_2 is 4 instead of 8.

3.2 The expected output of a layout algorithm

Definition 8. A class layout is a tuple of class-indexed families:

$$((dnvboff_C), (vboff_C), (foff_C), (nvdsiz_C), (nvsiz_C), (dsiz_C), (siz_C), (nvalign_C), (align_C))_{C \in \mathcal{C}}$$

where, for each class C :

- $\text{dnvboff}_C : \mathcal{DNV}(C) \rightarrow \mathbb{N}$ is a function assigning to each non-virtual direct base of C an offset within a subobject of type C . This offset corresponds to the direct non-virtual subobject of type B within a subobject of type C .
- $\text{vboff}_C : \mathcal{V}(C) \rightarrow \mathbb{N}$ is a function assigning to each virtual base of C an offset within a full instance of C .
- $\text{foff}_C : \text{ScF}(C) \cup \text{StF}(C) \rightarrow \mathbb{N}$ is a function assigning to each (scalar or structure array) field of C an offset within a subobject of type C .
- nvdsize_C is the non-virtual data size of C
- dsize_C is the data size of C
- nvsiz_C is the non-virtual size of C
- size_C is the size of C as would be given by the `sizeof` operator
- nvalign_C is the non-virtual alignment of C
- align_C is the alignment of C

As the hierarchy is assumed to be well-founded, a class layout algorithm would be able to incrementally compute offsets for a given class C , assuming that for all classes $B \prec C$, offsets have already been computed.

Again, the Coq development does not use dependent types to represent a class layout. Instead, a class layout is a mapping $O \in \text{PTree.t Offsets.t}$ where `Offsets.t` is a record type having fields such as `dnvboff`, `vboff`, `foff`, `...`. To enforce the domains, we externally constrain such mappings o with *guard constraints*:

- To incrementally compute the layout of a class C , it is assumed that for all classes $B \prec C$, their layout data are defined:

$$\forall B \in \mathcal{C} : B \prec C \Rightarrow O(B) \neq \emptyset$$

- Conversely, we assume that all layout data are computed only for valid classes:

$$\forall B \in \mathbb{N}^* : O(B) \neq \emptyset \Rightarrow B \in \mathcal{C}$$

- For a class B such that $O(B) = \{o\}$ for some layout data $o \in \text{Offsets.t}$, the function dnvboff_B is represented by a mapping $o.\text{dnvboff} \in \text{PTree.t } \mathbb{N}^*$ instead of $\text{PTree.t } \mathcal{DNV}(B)$ to avoid dependent types. So, the values have to be externally constrained:

$$\forall A \in \mathbb{N}^* : o.\text{dnvboff}(A) \neq \emptyset \Leftrightarrow A \in \mathcal{DNV}(B)$$

We put those constraints along with the layout constraints of Section 3.3. Then, given such a class layout, it is possible to compute the offset of non-virtual paths within a subobject of some class C , and the offset of any inheritance path within a full instance of C .

Definition 9. *If l is a non-virtual path from some class C to some class A , then $\text{nvsoff}(l) \in \mathbb{N}$ shall denote its offset within a subobject of type C . It is computed as follows, by structural recursion on l :*

- $\text{nvsoff}(C :: \text{nil}) \stackrel{\text{def.}}{=} 0$ ($C :: \text{nil}$ is the trivial non-virtual path to the subobject itself).
- if $l = C :: B :: l'$, then $B \in \mathcal{DNV}(C)$ and $B :: l'$ is a non-virtual path from B to A , and $\text{nvsoff}(C :: B :: l') \stackrel{\text{def.}}{=} \text{dnvboff}_C(B) + \text{nvsoff}(B :: l')$.

In our Coq formalization, we represented nvsoff through a tail-recursive function nvsoff' with an additional accumulating argument:

$$\begin{aligned} \text{nvsoff}'(\text{accu}, C :: \text{nil}) &\stackrel{\text{def.}}{=} \text{accu} \\ \text{nvsoff}'(\text{accu}, C :: B :: l') &\stackrel{\text{def.}}{=} \text{nvsoff}'(\text{accu} + \text{dnvboff}_C(B), B :: l') \end{aligned}$$

and we proved a rewriting lemma to show

$$\text{nvsoff}'(\text{accu}_1, l) = \text{nvsoff}'(\text{accu}_2, l) + \text{accu}_1 - \text{accu}_2$$

for any non-virtual path l and any accumulators $\text{accu}_1, \text{accu}_2 \in \mathbb{Z}$. So:

$$\text{nvsoff}(l) = \text{nvsoff}'(0, l)$$

Definition 10. *If (h, l) is a path from C to A , then its offset $\text{soff}_C(h, l) \in \mathbb{N}$ within a full instance of C is computed as follows:*

- $\text{soff}_C(\text{Repeated}, l) \stackrel{\text{def.}}{=} \text{nvsoff}(l)$,
- $\text{soff}_C(\text{Shared}, l) \stackrel{\text{def.}}{=} \text{vboff}_C(B) + \text{nvsoff}(l)$, where $B \in \mathcal{V}(C)$ and l is a non-virtual path from B to A

The above definition can be simplified by the convention $\text{vboff}_C(C) \stackrel{\text{def.}}{=} 0$, thanks to the well-foundedness of the hierarchy: indeed, C is not a virtual base of itself, so $\text{vboff}_C(C)$ is theoretically undefined. Then, with this convention, we obtain for any path $(h, B :: l')$ from C to A :

$$\text{soff}_C(h, B :: l') = \text{vboff}_C(B) + \text{nvsoff}(B :: l')$$

regardless of h . Remember that Lemma 7 (p. 12) makes h depend on whether $B = C$ or not. This leads us to introduce the notion of *generalized virtual bases* of a class.

Definition 11. The generalized virtual bases of C are C and the virtual bases of C . This set, denoted $\tilde{\mathcal{V}}(C)$, shall be the domain of $v\text{boff}_C$:

$$\begin{aligned} \tilde{\mathcal{V}}(C) &\stackrel{\text{def.}}{=} \mathcal{V}(C) \cup \{C\} \\ v\text{boff}_C : \quad \tilde{\mathcal{V}}(C) &\rightarrow \mathbb{N} \\ &C \quad \mapsto 0 \\ B \in \mathcal{V}(C) &\mapsto v\text{boff}_C(B) \end{aligned}$$

Lemma 8. If l' is a non-virtual path from some class B to some class A , then casting a subobject from B to A through path l' is equivalent to adding the offset $nv\text{soff}(l')$ to the offset of the subobject. That is, for any path P from C to B :

$$\text{soff}_C(P @ (\text{Repeated}, l')) = \text{soff}_C(P) + nv\text{soff}(l')$$

This shows that non-virtual derived-to-base casting does not depend on the starting subobject. However, to make a derived-to-base static cast through virtual inheritance, it is necessary to know the offset of the virtual base from the *full* object, so a subobject has to know its offset from the full object. See for instance Chen [3] who mechanically formalizes a type system for C++ multiple inheritance to determine which arithmetic operations are necessary for casts.

Given those definitions for inheritance path offsets, we can compute the offset of an array of n' structures of type C' within an array of n structures of size C :

Definition 12. Let C, C' be two classes, and n, n' two integers. If l is an array path from $C[n]$ to $C'[n']$, then the offset $a\text{off}_C(l)$ of $C'[n']$ within $C[n]$ through the array path l is computed as follows:

•

$$a\text{off}_C(\text{nil}) \stackrel{\text{def.}}{=} 0$$

- if $l = (i, P, f) :: l'$, then i is the index in the array $C[n]$, so $0 \leq i < n$, and there is a class B such that P is an inheritance path from C to B , and f is a field of B declared as an array of q structures of type A for some q and A , so that l' is an array path from $A[q]$ to $C'[n']$. Then we define:

$$a\text{off}_C((i, P, f) :: l') \stackrel{\text{def.}}{=} i \cdot \text{size}_C + \text{soff}_C(P) + \text{foff}_B(f) + a\text{off}_A(l')$$

Definition 13 (Relative pointer). Let $p = (\mathcal{P}, i, P)$ be a generalized subobject (or relative pointer) from an array of some type $C[n]$, where \mathcal{P} is an array path from $C[n]$ to $C'[n']$ for some class C' . Then, the offset of p within $C[n]$ is denoted:

$$\text{off}(p) \stackrel{\text{def.}}{=} a\text{off}_C(\mathcal{P}) + i \cdot \text{size}_{C'} + \text{soff}_{C'}(P)$$

3.3 Soundness

In our formalism, rather than imposing a fixed object layout as for C structures in CompCert [2], we propose some constraints on object layout, which layout algorithms must abide by. Then we prove that those constraints make two distinct scalar fields disjoint and two pointers to different subobjects of the same static type point to different memory locations.

Remark: Constraint equations and inequations are numbered following the same scheme as in the POPL 2011 draft submission [6]. They may appear in a different order, though.

Consider a class C .

3.3.1 Total size

The total size size_C of a class C is considered as an upper bound on the sizes of all its components. In the same way, the total non-virtual size nvsize_C of a class C is an upper bound on the sizes of all its non-virtual components (the non-virtual sizes of its direct non-virtual bases, and the sizes of its fields).

Hypothesis 3. *The non-virtual size of C is an upper bound on the non-virtual sizes of all its direct non-virtual bases:*

$$\forall B \in \mathcal{DNV}(C) : \text{dnvboff}_C(B) + \text{nvsize}_B \leq \text{nvsize}_C \quad (\text{C1})$$

Lemma 9. *If l is a non-virtual path from C to some class A , then the non-virtual part of A is included in the non-virtual part of C :*

$$\text{nvsoff}(l) + \text{nvsize}_A \leq \text{nvsize}_C$$

Proof. By induction on the length of l .

If l is the trivial path $C :: \text{nil}$, then $A = C$, $\text{nvsoff}(l) = 0$ and the inequality is an equality.

If $l = C :: B :: l'$, then $B \in \mathcal{DNV}(C)$ and $B :: l'$ is a non-virtual path from B to A , and we have $\text{nvsoff}(C :: B :: l') = \text{dnvboff}_C(B) + \text{nvsoff}(B :: l')$. By induction hypothesis, we have $\text{nvsoff}(B :: l') + \text{nvsize}_A \leq \text{nvsize}_B$. By the above hypothesis, we have $\text{dnvboff}_C(B) + \text{nvsize}_B \leq \text{nvsize}_C$, which concludes. \square

Hypothesis 4. *The total size of C is an upper bound on the non-virtual sizes of all its generalized virtual bases (including C itself):*

$$\forall B \in \tilde{\mathcal{V}}(C) : \text{vboff}_C(B) + \text{nvsize}_B \leq \text{size}_C \quad (\text{C3})$$

Lemma 10. *If (h, l) is an inheritance path (= base class subobject) from C to some class A , then the non-virtual part of A is included in C :*

$$\text{soff}_C(h, l) + \text{nvsize}_A \leq \text{size}_C$$

Proof. Recall that $l = B :: l'$ for some class $B \in \tilde{\mathcal{V}}(C)$ and some l' such that $B :: l'$ is a non-virtual path from B to A . We then have $\text{soff}_C(h, l) = \text{vboff}_C(B) + \text{nvsoff}(B :: l')$. (C3, p. 20) gives us $\text{vboff}_C(B) + \text{nvsizes}_B \leq \text{size}_C$. Lemma 9 (p. 19) gives us $\text{nvsoff}(B :: l') + \text{nvsizes}_A \leq \text{nvsizes}_B$, which concludes. \square

The non-virtual part of a class C contains all fields defined in C . To express this condition, we must define the *size* of a field.

Notation 4. We assume that there exists a function $\text{scsize} : \mathcal{A} \cup \mathcal{C} \rightarrow \mathbb{N}^*$ computing the size of a scalar value (atom, or pointer to C for some class C).

The function scsize should not depend on the layout algorithm. For most compilers, the size of a value of type pointer to T does not depend on T . But such assumption is not relevant in our formalization.

Definition 14. The size $\text{fsize}(f)$ of a field f is defined as follows:

- $\text{fsize}(f, t) \stackrel{\text{def.}}{=} \text{scsize}(t)$ for a scalar field
- $\text{fsize}(f, B, n) \stackrel{\text{def.}}{=} n \cdot \text{size}_B$ for a structure array field

Hypothesis 5. The total non-virtual size of C is an upper bound on the sizes of all its fields:

$$\forall F \in \mathcal{F}(C) : \text{foff}_C(F) + \text{fsize}(F) \leq \text{nvsizes}_C \quad (\text{C2})$$

Lemma 11. If l is an array path from $C[n]$ to $C'[n']$, then the designated array $C'[n']$ is “included” in $C[n]$:

$$\text{aoff}_C(l) + n' \cdot \text{size}_{C'} \leq n \cdot \text{size}_C$$

Proof. By structural induction on l .

- If $l = \text{nil}$, then $C' = C$ and $n' = n$ and $\text{aoff}_C(l) = 0$, so the inequality is an equality.
- Otherwise, $l = (i, (h, p), f) :: l'$ where $0 \leq i < n$, (h, p) is an inheritance path from C to some class B , and f is a structure array of some type $A[m]$ defined in B . So, we have $\text{aoff}_C(l) = i \cdot \text{size}_C + \text{soff}_C(h, p) + \text{foff}_B(f) + \text{aoff}_A(l')$, and :
 - by induction hypothesis, $\text{aoff}_A(l') + n' \cdot \text{size}_{C'} \leq m \cdot \text{size}_A = \text{fsize}(f)$ by definition
 - $\text{foff}_B(f) + \text{fsize}(f) \leq \text{nvsizes}_B$ by (C2, p. 20)
 - $\text{soff}_C(h, p) + \text{nvsizes}_B \leq \text{size}_C$ by Lemma 10 (p. 20)
 - $i + 1 \leq n$ by hypothesis

which concludes. □

Corollary 12. *If p is a relative pointer (= generalized subobject) from $C[n]$ to some class A , then the non-virtual part of A is included in $C[n]$:*

$$\text{off}_C(p) + \text{nvsiz}_A \leq n \cdot \text{size}_C$$

Proof. Immediately follows from Lemmata 10 (p. 20) and 11 (p. 21). □

3.3.2 Alignment

The non-virtual alignment of a class C is a boundary used to align C as a non-virtual base of another class. So, it has to be a multiple of the alignments of all the non-virtual components of C : the alignments of fields defined in C , and the non-virtual alignments of non-virtual bases of C .

So we must first define field alignments.

Notation 5. *We assume that there exists a function $\text{scalign} : \mathcal{A} \cup \mathcal{C} \rightarrow \mathbb{N}^*$ computing the alignment of a scalar value (atom, or pointer to B for some class B).*

As with scsize , the function scalign should not depend on the layout algorithm.

Definition 15. *The alignment $\text{falign}(f)$ of a field f is defined as follows:*

- $\text{falign}(f, t) \stackrel{\text{def.}}{=} \text{scalign}(t)$ for a scalar field
- $\text{falign}(f, B, n) \stackrel{\text{def.}}{=} \text{align}_B$ for a structure array field

Indeed, a structure array field contains “full” instances of some class B , so the field must be aligned to the whole alignment of B , not only the non-virtual alignment of B .

Hypothesis 6. *The non-virtual alignment of a class C is a multiple of the alignments of its fields:*

$$\forall F \in \mathcal{F}(C) : (\text{falign}(F) \mid \text{nvalign}_C) \tag{C14b}$$

Hypothesis 7. *Any field f of C must be laid out with respect to its alignment:*

$$\forall F \in \mathcal{F}(C) : (\text{falign}(F) \mid \text{foff}_C(f)) \tag{C14a}$$

Hypothesis 8. *The non-virtual alignment of a class C is a multiple of the non-virtual alignments of its direct non-virtual bases:*

$$\forall B \in \mathcal{DNV}(C) : (\text{nvalign}_B \mid \text{nvalign}_C) \tag{C15b}$$

Hypothesis 9. Any non-virtual base B of C must be laid out with respect to its non-virtual alignment:

$$\forall B \in \mathcal{DNV}(C) : (n\text{align}_B \mid d\text{nvboff}_C(B)) \quad (\text{C15a})$$

Finally, the alignment of a class C synthesizes the non-virtual alignments of C and its virtual bases.

Hypothesis 10. The alignment of a class C is a multiple of the non-virtual alignments of its generalized virtual bases (including C itself):

$$\forall B \in \tilde{\mathcal{V}}(C) : (n\text{align}_B \mid \text{align}_C) \quad (\text{C17b})$$

Hypothesis 11. Any virtual base B of C must be laid out with respect to its non-virtual alignment:

$$\forall B \in \mathcal{V}(C) : (n\text{align}_B \mid \text{vboff}_C(B)) \quad (\text{C17a})$$

(C17a, p. 22) is also true for $B = C$, as $\text{vboff}_C(C) \stackrel{\text{def.}}{=} 0$.

Then, all those constraints allow to show that the access to a field is correctly aligned.

Lemma 13. If l is a non-virtual path from some class C to some class A , then the access to A through l is correctly aligned:

$$(n\text{align}_A \mid n\text{align}_C) \quad (i)$$

$$(n\text{align}_A \mid \text{nvsoff}(l)) \quad (ii)$$

Proof. By structural induction on l .

- If $l = \text{nil}$, then $C = A$ and $\text{nvsoff}(l) = 0$, which trivially concludes.
- Otherwise, $l = B :: l'$ with $B \in \mathcal{DNV}(C)$ and $\text{nvsoff}(l) = d\text{nvboff}_C(B) + \text{nvsoff}(l')$. So:
 - $(n\text{align}_B \mid n\text{align}_C)$ by (C14a, p. 22) and $(n\text{align}_A \mid n\text{align}_B)$ by induction hypothesis, thus (i) by transitivity.
 - $(n\text{align}_B \mid d\text{nvboff}_C(B))$ by (C14b, p. 22), and $(n\text{align}_A \mid n\text{align}_B)$ by induction hypothesis, and $(n\text{align}_A \mid \text{nvsoff}(l'))$ by induction hypothesis, thus (ii) by transitivity and compatibility of $(. \mid .)$ with addition.

□

Lemma 14. If (h, l) is an inheritance path (= base class subobject) from some class C to some class A , then the access to A through (h, l) is correctly aligned:

$$(n\text{align}_A \mid \text{align}_C) \quad (i)$$

$$(n\text{align}_A \mid \text{soff}_C(h, l)) \quad (ii)$$

Proof. Recall that $l = B :: l'$ for some class $B \in \tilde{\mathcal{V}}(C)$ and some l' such that $B :: l'$ is a non-virtual path from B to A . We then have $\text{soff}_C(h, l) = \text{vboff}_C(B) + \text{nvsoff}(B :: l')$.

- (C17a, p. 22) gives us $(\text{nvalign}_B \mid \text{align}_C)$; part (i) of Lemma 13 (p. 22) gives us $(\text{nvalign}_A \mid \text{nvalign}_B)$, thus (i) by transitivity.
- (C17b, p. 22) gives us $(\text{nvalign}_B \mid \text{vboff}_C(B))$; Lemma 13 (p. 22) gives us $(\text{nvalign}_A \mid \text{nvalign}_B)$ and $(\text{nvalign}_A \mid \text{nvsoff}(l'))$, thus (ii) by transitivity and compatibility of $(. \mid .)$ with addition.

□

If C is a class and c is an array of several structures of type C , then access to the second structure item of c will be realized at low-level through adding an offset of $\text{sizeof}(C)$ bytes to a pointer to c . So, the following constraint is necessary to ensure the correct alignment of such an access:

Hypothesis 12.

$$(\text{align}_C \mid \text{size}_C) \tag{C18}$$

This condition will ensure the:

Lemma 15. *If l is an array path from $C[n]$ to $C'[n']$, then the designated array $C'[n']$ is correctly aligned in $C[n]$:*

$$(\text{align}_{C'} \mid \text{align}_C) \tag{i}$$

$$(\text{align}_{C'} \mid \text{aoff}_C(l)) \tag{ii}$$

Proof. By structural induction on l .

- If $l = \text{nil}$, then $C' = C$ and $\text{aoff}_C(l) = 0$, which trivially concludes.
- Otherwise, $l = (i, (h, p), f) :: l'$ where $0 \leq i < n$, (h, p) is an inheritance path from C to some class B , and f is a structure array of some type $A[m]$ defined in B , such that $\text{falign}(f) = \text{align}_A$. So, we have $\text{aoff}_C(l) = i \cdot \text{size}_C + \text{soff}_C(h, p) + \text{foff}_B(f)$, and:
 - $(\text{align}_{C'} \mid \text{align}_A)$ by induction hypothesis, and $(\text{falign}(f) \mid \text{nvalign}_B)$

$$\begin{array}{c} \parallel \\ \text{align}_A \end{array}$$
by (C15b, p. 22), and $(\text{nvalign}_B \mid \text{align}_C)$ by part (i) of Lemma 14 (p. 23), thus (i) by transitivity
 - thanks to (i) and (C18, p. 23), we have $(\text{align}_{C'} \mid i \cdot \text{size}_C)$. Moreover, by induction hypothesis, $(\text{align}_{C'} \mid \text{align}_A)$, and $(\text{align}_A \mid \text{nvalign}_B)$ as above, and $(\text{nvalign}_B \mid \text{soff}_C(h, p))$ by part (ii) of Lemma 14 (p. 23), and $(\text{align}_A \mid \text{foff}_B(f))$ by (C15a, p. 22), and $(\text{align}'_C \mid \text{aoff}_A(l'))$ by induction hypothesis, thus (ii).

□

Corollary 16. *If p is a relative pointer (= generalized subobject) from $C[n]$ to some class A , then the access to A through p is correctly aligned:*

$$(nvalign_A \mid align_C)$$

$$(nvalign_A \mid off_C(p))$$

Proof. Immediately follows from Lemmata 14 (p. 23) and 15 (p. 24). □

This lemma stresses out the difference of use between the alignment and the non-virtual alignment of a class: whereas full instances are aligned to the full alignment of their class, by contrast, generalized subobjects are aligned to the non-virtual alignment of their static type.

This finally allows to show the:

Theorem 1. *If p is a generalized subobject of static type A from a structure array of C , and if (f, t) is a scalar field of A of scalar type $t \in \mathcal{A} \cup \mathcal{C}$, then the access to this field is correctly aligned:*

$$(scalalign(t) \mid align_C)$$

$$(scalalign(t) \mid off_C(p))$$

Proof. Immediately follows from Corollary 16 (p. 24) and conditions (C15a, p. 22) and (C15b, p. 22). □

3.3.3 Data size

We aim at stating constraints to forbid the overlapping of different scalar data. So, we shall only consider here *non-empty* components, in the sense that roughly speaking, a component is non-empty as soon as it contains some scalar data.

Then, to express that different scalar data do not overlap, we shall not use the sizes, but the *data sizes* of the relevant components. In this section, we shall show that data sizes are not relevant for empty components.

Recall that the actual definition for an empty class is left to the layout algorithm. Based on this definition, we may introduce the notion of *empty field*, which is a structure field of an empty class type, regardless of the number of its array elements:

Notation 6. *The set of empty fields of C is the set of those fields defined in class C that are array structures of type B for some empty class B :*

$$EF(C) \stackrel{\text{def.}}{=} \{(f, B, n) \in StF(C) \mid B \text{ is empty}\}$$

The set of non-empty fields of C is:

$$NEF(C) \stackrel{\text{def.}}{=} ScF(C) \cup StF(C) \setminus EF(C)$$

Then, we may define the notion of the *data size* of a field. For a scalar field, the data size of a field is equal to its size. But for a structure array field, there are two cases:

- if the field has only one structure element of some type B , then the data size of such a field is equal to the data size of the class B .
- otherwise, it is important to note that the data of a field is considered to be a *contiguous interval*. So, any padding between two array elements is lost. However, nothing prevents us from reusing the padding of the last element of the array.

For this reason, we define the *data size* of a non-empty field as follows.

Definition 16. *The data size of a scalar field (f, t) is:*

$$fsize(f, t) \stackrel{\text{def.}}{=} scsize(t)$$

The data size of a non-empty structure array field (f, b, n) is:

$$fsize(f, b, n) \stackrel{\text{def.}}{=} (n - 1) \cdot size_B + dsize_B$$

Then, in the same way as for the whole sizes, we introduce the notions of *non-virtual data size* (for the non-virtual part of a class) and *data size* (for the whole class), as bounds on the sizes of the class components.

Notation 7. *In the same way as for fields, we shall use the following notations:*

- $NEDNV(C)$ for the set of non-empty direct non-virtual bases of C
- $EDNV(C)$ for the set of empty direct non-virtual bases of C
- $NEV(C)$ for the set of non-empty virtual bases of C
- $EV(C)$ for the set of empty virtual bases of C
- $NE\tilde{V}(C)$ for the set of non-empty generalized virtual bases of C
- $E\tilde{V}(C)$ for the set of empty generalized virtual bases of C

Then, the non-virtual data of a class C is divided into two parts, each being a *contiguous interval*:

- the non-virtual data of the non-empty direct non-virtual bases of C
- the data of the fields defined in C

Hypothesis 13. *There exists a boundary $fboundary_C \in \mathbb{N}$ between the non-virtual data for the direct non-virtual bases of C and the non-empty fields of C :*

$$\forall B \in NEDNV(C) : dnvboff_C(B) + nvsize_B \leq fboundary_C \quad (C7)$$

$$\forall F \in NEF(C) : fboundary_C \leq foff_C(F) \quad (C8)$$

That hypothesis ensures that the two parts are disjoint.

In our formalization, such boundary is included in the definition of a class layout (i.e. $fboundary_C$ is a value explicitly computed by the layout algorithm).

Hypothesis 14. *The non-virtual data size of class C is an upper bound of the sizes of the two parts of the non-virtual data of C :*

$$fboundary_C \leq nvsize_C \quad (C11)$$

$$\forall F \in NEF(C) : foff_C(F) + fsize(F) \leq nvsize_C \quad (C10)$$

(C11, p. 26) is necessary if there are no non-empty fields defined in C . Otherwise, (C11, p. 26) is entailed by (C10, p. 26).

Lemma 17. *If l is a non-virtual path from C to some non-empty class A , then the non-virtual part of A is included in the non-virtual part of C :*

$$nvsoff(l) + nvsize_A \leq nvsize_C$$

Proof. Same shape as the proof for Lemma 9 (p. 19), additionally using (C11, p. 26). Additionally to A , all considered classes are not empty, as they have a non-empty base A . \square

Hypothesis 15. *The total data size of C is an upper bound on the non-virtual data sizes of all its non-empty generalized virtual bases (including C itself, if non-empty):*

$$\forall B \in NE\tilde{V}(C) : vboff_C(B) + nvsize_B \leq dsize_C \quad (C13)$$

Lemma 18. *If (h, l) is an inheritance path (= base class subobject) from C to some class A , then the non-virtual data of A is included in C :*

$$soff_C(h, l) + nvsize_A \leq dsize_C$$

Proof. Same shape as the proof for Lemma 10 (p. 20). \square

Recall that the *data* of a structure array of n cells of type C consisting of a *contiguous interval* embedding the total sizes of the first $n - 1$ cells and the only data of the last cell:

$$(n - 1) \cdot size_C + dsize_C$$

Lemma 19. *If $p = (l, i, P)$ is a relative pointer (= generalized subobject) from $C[n]$ to some non-empty class A , then the non-virtual data of A is included in the data of $C[n]$:*

$$\text{off}_C(p) + \text{nvdsiz}_A \leq (n - 1) \cdot \text{size}_C + \text{dsiz}_C$$

Proof. By structural induction on l .

- If $l = \text{nil}$, then $i \leq (n - 1)$ and $\text{off}_C(p) = i \cdot \text{size}_C + \text{soff}_C(P)$; moreover, P is an inheritance path from C to A , so by Lemma 18 (p. 27), $\text{soff}_C(P) + \text{nvdsiz}_A \leq \text{dsiz}_C$, which concludes.
- Otherwise, $l = (i', (h', p'), f') :: l'$ where $0 \leq i' < n$, (h', p') is an inheritance path from C to some class B , and f' is a structure array of some type $C'[n']$ defined in B , such that (l', i, P) is a relative pointer from $C'[n']$ to A , and $\text{off}_C(p) = i \cdot \text{size}_C + \text{soff}_C(h, p) + \text{foff}_B(f') + \text{off}_{C'}(l', i, P)$, and :
 - by induction hypothesis, $\text{off}_{C'}(l', i, P) + \text{nvdsiz}_A \leq (n' - 1) \cdot \text{size}_{C'} + \text{dsiz}_{C'} = \text{fsiz}(f')$ by definition
 - $\text{foff}_B(f) + \text{fsiz}(f) \leq \text{nvdsiz}_B$ by (C11, p. 26) (legal because, as A is not empty, neither is B)
 - $\text{soff}_C(h, p) + \text{nvdsiz}_B \leq \text{dsiz}_C$ by Lemma 18 (p. 27)
 - $i' \leq n - 1$ by hypothesis

which concludes. □

3.3.4 Non-overlapping of data

To express that two fields F_1 and F_2 defined in some class C should not overlap, we could expect their full size intervals to be disjoint:

$$[\text{foff}_C(F_1), \text{foff}_C(F_1) + \text{fsiz}(F_1)) \perp [\text{foff}_C(F_2), \text{foff}_C(F_2) + \text{fsiz}(F_2))$$

This condition is actually enforced by the Itanium C++ ABI, as seen in Section 4.1. However, it impedes the reusing of tail padding, even alignment tail padding. For instance, consider the following structure:

```

struct Z {};
struct A: Z {
    int i;
    char ca;
};
struct B {
    A a;
    char cb;
}

```

Most compilers (such as GCC) lay out a and cb completely disjointly within B , laying out cb at offset 8 (on Intel x86 32-bit platforms). This incurs the loss of unused space between the end of $a.ca$ and the beginning of cb : 3 bytes for alignment padding.

Our formalization allows reusing this padding, by defining *data sizes* distinct from sizes, trying to exclude padding as most as possible. Then, we introduce constraints to ensure that *data* intervals be disjoint.

In our example, whereas the size of a is 8, its data size would be only 5 thanks to our formalization, such that cb would be laid out at offset 5 instead of 8. Thus, we have to use the notion of *data size* rather than size, to express that two fields of the same class do not overlap:

Hypothesis 16. *The data of two non-empty fields defined in the same class do not overlap:*

$\forall F_1, F_2 \in NE\mathcal{F}(C) :$

$$\begin{aligned} & [foff_C(F_1), foff_C(F_1) + fsize(F_1)) \\ \perp & [foff_C(F_2), foff_C(F_2) + fsize(F_2)) \end{aligned} \quad (C9)$$

Similarly, the non-virtual data of two direct non-virtual bases of C do not overlap:

Hypothesis 17. $\forall B_1, B_2 \in DN\mathcal{V}(C) :$

$$\begin{aligned} & [dnvboff_C(B_1), dnvboff_C(B_1) + nvsize_{B_1}) \\ \perp & [dnvboff_C(B_2), dnvboff_C(B_2) + nvsize_{B_2}) \end{aligned} \quad (C6)$$

However, non-virtual sizes may overlap. Consider for instance the following hierarchy:

```

struct A      {} ;
struct B      { int i } ;
struct C : A, B {} ;
```

Then, the following layout (actually given by GNU G++) complies to the above constraints:

$$\begin{aligned} nvsize_A &= 0 \\ nsize_A &= 1 \\ foff_B(i, \text{int}) &= 0 \\ nvsize_B &= \text{scsize}(\text{int}) \\ nsize_B &= \text{scsize}(\text{int}) \\ dnvboff_C(A) &= 0 \\ dnvboff_C(B) &= 0 \\ nvsize_C &= \text{scsize}(\text{int}) \\ nsize_C &= \text{scsize}(\text{int}) \end{aligned}$$

Here, the non-virtual sizes of A and B overlap in C , but this does not impede the access to field i of B , as A is empty.

The actual purpose of distinguishing sizes from data sizes is that a pointer to an empty class can point outside of the actual data (accessible non-empty fields), but must still point inside the whole size of the object (so as to prevent it from pointing into another object that would be created independently). However, we shall see further in more detail tighter conditions to prevent two pointers to different subobjects of static type A (where A is an empty class) from pointing to the same memory location.

Lemma 20. *If l_1, l_2 are two distinct non-virtual paths from some class C to some non-empty classes B_1, B_2 , then their field data zones are disjoint:*

$$\perp \begin{array}{l} [nvsoff(l_1) + fboundary_{B_1}, nvsoff(l_1) + nvdsiz_{B_1}) \\ [nvsoff(l_2) + fboundary_{B_2}, nvsoff(l_2) + nvdsiz_{B_2}) \end{array}$$

Proof. For symmetry reasons, we may assume $\text{length}(l_1) \leq \text{length}(l_2)$. Then, there are two cases:

- either there is a non-virtual non-trivial path $B_1 :: A :: l'$ from B_1 to B_2 such that $l_2 = l_1 @_{\text{Repeated}}(B_1 :: A :: l')$. So A is a non-virtual direct base of B_1 , and actually the non-virtual data of A (which includes the non-virtual data of B_2 , in particular field F_2) is disjoint from the field data of B_1 (in particular field F_1) thanks to the boundary $fboundary_{B_1}$.
- or there are a class A and two distinct non-virtual direct bases A_1 and A_2 of A , and a list L such that for each $i \in \{1, 2\}$, $L+A :: A_i :: \text{nil}$ is a non-virtual path from C to A_i and $A_i :: l'_i$ is a non-virtual path from A_i to B_i for some l'_i . The non-virtual direct bases A_1 and A_2 of A have disjoint non-virtual data. The non-virtual data of A_i includes the non-virtual data of B_i , and in particular field F_i , so those fields are disjoint.

□

Recall the convention $\mathbf{vboff}_C(C) \stackrel{\text{def.}}{=} 0$. This allows to roughly say that the non-virtual data of all non-empty generalized virtual bases of C (with C considered as a generalized virtual base of itself) are laid out one after another, in such a way that they are disjoint:

Hypothesis 18. $\forall B_1, B_2 \in NE\tilde{\mathcal{V}}(C) : \text{if } B_1, B_2 :$

$$\perp \begin{array}{l} [vboff_C(B_1), vboff_C(B_1) + nvdsiz_{B_1}) \\ [vboff_C(B_2), vboff_C(B_2) + nvdsiz_{B_2}) \end{array} \quad (\text{C12})$$

It follows immediately that:

Lemma 21. *If P_1, P_2 are distinct inheritance paths from C to some B_1, B_2 , then their field data zones are disjoint:*

$$\perp \begin{array}{l} \left[\text{soff}_C(P_1) + \text{fboundary}_{B_1}, \text{soff}_C(P_1) + \text{nvdsiz}_{B_1} \right) \\ \left[\text{soff}_C(P_2) + \text{fboundary}_{B_2}, \text{soff}_C(P_2) + \text{nvdsiz}_{B_2} \right) \end{array}$$

Proof. Let $P_i = (h_i, A_i :: l_i)$ for each i . If $A_1 = A_2$, then Lemma 20 (p. 29) about non-virtual paths can be reused. Otherwise, we know by (C12, p. 30) that the non-virtual data of A_i are disjoint, then Lemma 17 (p. 27) and (C11, p. 26) conclude. \square

The conditions given so far are enough to show that two different scalar fields reachable from two base class subobjects of an object are disjoint. However, when it comes to traversing structure array fields, we must show that two fields reachable from two different cells of the same array are disjoint. So far there is no condition ensuring such a property. Indeed, there is yet no link between the data size and the size of a class: we have to explicitly constrain that:

Hypothesis 19. *For any class C :*

$$\text{dsiz}_C \leq \text{siz}_C \tag{C4}$$

Notation 8. *The length of a pointer $p = (\mathcal{P}, j, P)$, denoted $\text{plength}(p)$, is the length of its array path \mathcal{P} :*

$$\text{plength}(p) \stackrel{\text{def.}}{=} \text{length}(\mathcal{P})$$

Theorem 2. *If p_i are two generalized subobjects of static type B_i within a structure array of type C , and if F_i are two **scalar** fields defined in class B_i , such that $(p_1, F_1) \neq (p_2, F_2)$, then those two fields are disjoint.*

Proof. For symmetry reasons, we may assume $\text{plength}(p_1) \leq \text{plength}(p_2)$. Reason by induction on $\text{plength}(p_1)$. Then, we shall introduce an alternate representation for p_i :

Definition 17 (Alternate representation of generalized subobjects (= alternate relative pointers)). *For any generalized subobject (= relative pointer) $p = (\mathcal{P}, j, P)$ from a structure array of type C to some class B , we can find $j' \in \mathbb{Z}$, an inheritance path P' from C to some class B' , and an option Φ such that:*

- either $\Phi = \emptyset$ and $p = (\text{nil}, j', P')$
- or there is a structure array field F' and a relative pointer $p' =$ for some array path \mathcal{P}' such that $\Phi = \{(F', (\mathcal{P}', j, P))\}$ and $\mathcal{P} = (j', P', F') :: \mathcal{P}'$ (so that $\text{plength}(\mathcal{P}', j, P) = \text{plength}(p) - 1$ to allow induction).

(j', P', Φ) is called the alternate representation of p , and we note:

$$(\mathcal{P}, j, P) \propto (j', P', \Phi)$$

Lemma 22. *The alternate representation of a relative pointer p is unique:*

$$p \propto \mathfrak{p}_1 \wedge p \propto \mathfrak{p}_2 \Rightarrow \mathfrak{p}_1 = \mathfrak{p}_2$$

This alternate representation is intended for proofs about concrete object layout, by contrast to the “regular” Definition 6 (p. 9) which is designed for the abstract high-level semantics of C++ (especially to model casts and other dynamic operations such as virtual method dispatch).

This alternate representation is trivially compatible with offsets:

Lemma 23. *Let $p \propto (j', P', \Phi)$ be a relative pointer from C of static type B .*

- if $\Phi = \emptyset$, then:

$$\text{off}_C(p) = j' \cdot \text{size}_C + \text{soff}_C(P')$$

- otherwise, P' is an inheritance path from C to some class A and $\Phi = \{(F', p')\}$ for some structure field F' of type B' defined in A and for some relative pointer p' from B of static type B , and:

$$\text{off}_C(p) = j' \cdot \text{size}_C + \text{soff}_C(P') + \text{foff}_A(F') + \text{off}_{B'}(p')$$

In particular, the following interesting lemma holds:

Lemma 24. *If $p \propto (j', P', \Phi)$ is a relative pointer from C of static type B , such that P' is an inheritance path from C to some class C' , then the field data of p is included in the field data of C' :*

$$\begin{aligned} & [\text{off}_C(p) + \text{fboundary}_B, \text{off}_C(p) + \text{nvdsize}_B) \\ \subseteq & [j' \cdot \text{size}_C + \text{soff}_C(P') + \text{fboundary}_{C'}, j' \cdot \text{size}_C + \text{soff}_C(P') + \text{nvdsize}_{C'}] \end{aligned}$$

Then, for each i , let $p_i \propto (j'_i, P'_i, \Phi_i)$. There are several cases:

- if $j'_1 \neq j'_2$, then we have two disjoint cells of an array of structures of some type C . Thanks to (C4, p. 31), their data are also disjoint, which concludes.
- $j'_1 = j'_2$ and $P'_1 \neq P'_2$: then, thanks to the above lemma, the theorem for virtual inheritance directly applies.
- Now assume $j'_1 = j'_2$, $P'_1 = P'_2$.
 - If $\Phi_1 = \Phi_2$, then $F_1 \neq F_2$ are distinct fields of the same subobject, so by (C9, p. 28) they are disjoint.

- Otherwise, we can assume $\Phi_2 = \{F'_2, p'_2\}$ (because $\text{plength}(p_1) \leq \text{plength}(p_2)$, so that if $\Phi_2 = \emptyset$, then $\Phi_1 = \emptyset$).
 - * if $\Phi_1 = \emptyset$, then F_1 and F'_2 are distinct fields (because F_1 is scalar whereas F'_2 is a structure field) of the same class, so their data are disjoint.
 - * Otherwise, $\Phi_1 = \{F'_1, p'_1\}$.
 - if $F'_1 \neq F'_2$, then they are two distinct fields of the same class, so their data are disjoint.
 - Otherwise, we may use the induction hypothesis.

□

Our layout constraints allow to produce smarter layouts than GNU G++. Indeed, fields are laid out not by their whole sizes, but only by their data sizes. This allows fields to be stored within the end of a structure array field. Consider for instance:

```

struct A           { };
struct B1: A       { };
struct B2: A       { };
struct C: B1, B2   { char i; };
struct D           { C c; char j; };

```

Assuming $\text{sizeof}(\text{char}) = 1$, GNU G++ gives $\text{sizeof}(C) = 2$, as 2 different offsets to subobjects of type A have to be allotted.

But GNU G++ would give $\text{sizeof}(D) = 3$ by laying out j beyond the size of field c , so at offset 2, whereas our formalization allows j to be laid out at offset 1, that is just after the end of field i of c , not waiting for the actual end of field c . Such an optimization is proposed by [5].

In other words, our formalization generalizes empty base offsets to empty fields, thanks to the fact that the data size of a field of type $C[n]$ is $(n - 1) \cdot \text{size}_C + \text{dsz}_C$ instead of $n \cdot \text{size}_C$ as prescribed by Itanium C++ ABI [1].

3.3.5 Dynamic type data

If C is dynamic, then there are two cases:

- if C has a direct non-virtual base A that is dynamic, then this base may be chosen as a *primary base* that will have offset 0 within the non-virtual part of C .
- otherwise, some space must be explicitly allocated at the beginning of C to store the dynamic type data.

This indeed allows any dynamic class C to necessarily have some space at the beginning of C to store the dynamic type data. If C has a primary base A , then C and A will share their dynamic type data.

Hypothesis 20. *The size of dynamic type data is denoted $dtdsize$. It is positive and it does not depend on any class.*

In practice, for such compiler as GNU C++, $dtdsize = \text{sizeof}(\text{void}^*)$ to store a pointer to a virtual table. Other compilers like MSVC++ use $dtdsize = 2 * \text{sizeof}(\text{void}^*)$ to store additional data (a pointer to the table of virtual bases) [3].

The choice of the primary base is formalized as follows:

Hypothesis 21. *There exists $pbase_C \in \mathfrak{B}(\mathcal{DNV}(C))$ with cardinality 0 or 1 (that is, $pbase_C \in \text{option } \mathcal{DNV}(C)$), such that if $pbase_C = \{A\}$, then C and A are dynamic and:*

$$dnvboff_C(A) = 0 \quad (\text{C21})$$

In this case, A is called the primary base of C .

A class C having a dynamic non-virtual direct base does not necessarily have a primary class: the layout algorithm is not forced to choose any. However, as regards performance, such an algorithm could yield under-optimized layouts, with C having its own dynamic type data disjoint from all of its bases'.

We have to prove non-overlapping lemmas for dynamic type data. Indeed, we must ensure that:

- whenever a field is written to, the dynamic type data of dynamic sub-objects are not altered
- whenever a subobject is being initialized, the dynamic type data of other subobjects are not altered

Disjointness of fields and dynamic type data

Hypothesis 22. *The dynamic type data of a dynamic class C is disjoint from the field data of C :*

$$dtdsize \leq fboundary_C \quad (\text{C19})$$

This hypothesis along with (C11, p. 26) allows to show that the dynamic type data of a class is included in its non-virtual data.

The following hypothesis makes C 's dynamic type data disjoint from any of its non-empty non-virtual direct bases if C has no primary base:

Hypothesis 23. *If $pbase_C = \emptyset$, then for any non-empty direct non-virtual base B of a dynamic class C :*

$$dtdsize \leq dnvboff_C(B) \quad (\text{C20})$$

Lemma 25. *If l is a non-virtual path from some dynamic class C to some non-empty class B , then the dynamic type data of C is disjoint from the field data of B . More precisely:*

$$dtdsize \leq nvsoff(l) + fboundary_B$$

Proof. By induction on l . If $l = C :: \text{nil}$ is the trivial path, then $B = C$ and (C19, p. 34) concludes. Otherwise, if $l = C :: B' :: l'$, then B' is a non-virtual direct base of C . There are two cases:

- If B' is the primary base of C , having offset 0 within C , then B' is dynamic and we use the induction hypothesis.
- Otherwise, it suffices to show that B' itself starts at offset at least dtdsize within C , which would conclude. There are two cases:
 - If C has no primary base, then (C20, p. 34) concludes.
 - Otherwise, if P is the primary base of C , then the data of P and B' are disjoint. However, we need an additional hypothesis to conclude:

Hypothesis 24. *Any non-empty class A has non-null non-virtual data:*

$$0 < \text{nvdsiz}_A \tag{C22}$$

In practice, this hypothesis makes sense, but its proof depends on the layout algorithm and, in particular, the notions of empty class and dynamic class.

□

Lemma 26. *If l is a non-virtual path from some class C to some dynamic class B , then the dynamic type data of B is disjoint from the field data of C . More precisely:*

$$\text{nvsoff}(l) + \text{dtdsize} \leq \text{fboundary}_C$$

Proof. By case analysis on l . If $l = C :: \text{nil}$ is the trivial path, then $B = C$ and (C19, p. 34) concludes. Otherwise, if $l = C :: B' :: l'$, then B' is a non-virtual direct base of C , and:

- $\text{dtdsize} \leq \text{fboundary}_B$ by (C19, p. 34)
- $\text{fboundary}_B \leq \text{nvdsiz}_B$ by (C11, p. 26)
- $\text{nvsoff}(B' :: l) + \text{nvdsiz}_B \leq \text{nvdsiz}_{B'}$ by Lemma 17 (p. 27)
- $\text{dnvboff}_C(B') + \text{nvdsiz}_{B'} \leq \text{fboundary}_C$ by (C7, p. 26)

which concludes. □

Lemma 27. *If l_1 is a non-virtual path from C to some dynamic class B_1 , and if l_2 is a non-virtual path from C to some non-empty class B_2 , then the field data of B_2 is disjoint from the dynamic data of B_1 :*

$$\perp \left[\begin{array}{l} \text{nvsoff}(l_1), \text{nvsoff}(l_1) + \text{dtdsize} \\ \text{nvsoff}(l_2) + \text{fboundary}_{B_2}(F), \text{nvsoff}(l_2) + \text{nvdsiz}_{B_2} \end{array} \right]$$

Proof. There are three cases:

- if there is a non-virtual path p from B_1 to B_2 such that $l_2 = l_1 @_{\text{Repeated}} p$, then Lemma 25 (p. 34) concludes.
- if there is a non-virtual path p from B_2 to B_1 such that $l_1 = l_2 @_{\text{Repeated}} p$, then Lemma 26 (p. 35) concludes.
- Otherwise, there is a class A and a non-virtual path p from C to A and $A_1 \neq A_2 \in \mathcal{DNV}_A$ and non-virtual paths l'_1 from A_1 to B_1 and l'_2 from A_2 to B_2 such that $\forall i : l_i = p @_{\text{Repeated}} (A :: A_i :: \text{nil}) @_{\text{Repeated}} l'_i$. As A_1 and A_2 are two distinct direct non-virtual non-empty bases of A , their data zones are disjoint, which concludes.

□

Corollary 28. *If P_1 is an inheritance path from C to some dynamic class B_1 , and if P_2 is an inheritance path from C to some non-empty class B_2 , then the field data of B_2 is disjoint from the dynamic data of B_1 :*

$$\perp \begin{array}{l} [\text{soff}_C(P_1), \text{soff}_C(P_1) + \text{dtdsize}] \\ [\text{soff}_C(P_2) + \text{fboundary}_{B_2}(F), \text{soff}_C(P_2) + \text{nvdsiz}_{B_2}] \end{array}$$

Proof. Let $P_i = (h_i, B'_i :: l_i)$ for each i . Then, there are two cases:

- if $B'_1 = B'_2$, then the previous lemma immediately concludes.
- Otherwise, B'_1 and B'_2 are two distinct non-empty generalized virtual bases of C , so by (C12, p. 30) their data are disjoint, which concludes thanks to Lemma 18 (p. 27)

□

Theorem 3. *If p_1, p_2 are generalized subobjects of static type B_1, B_2 within an array of C , such that B_1 defines a scalar field F and B_2 is dynamic, then F is disjoint from the dynamic type data of B_1 :*

$$\perp \begin{array}{l} [\text{off}_C(p_1) + \text{foff}_{B_1}(F), \text{off}_C(P_1) + \text{foff}_{B_1}(F) + \text{fdsiz}(F)] \\ [\text{off}_C(p_2), \text{off}_C(p_2) + \text{dtdsize}] \end{array}$$

Proof. By induction on $\text{plength}(p_1)$. For each i , let $p_i \propto (j_i, P_i, \Phi_i)$. Then:

- If $j_1 \neq j_2$, then we have two different cells of the array of C , so thanks to (C4, p. 31) and Lemma 19 (p. 27), their data are disjoint.

Otherwise, let B'_i be the static type of P_i . By Lemma 24 (p. 32), F is included in the field data of B'_1 .

- If $\Phi_2 = \emptyset$, then Corollary 28 (p. 36) concludes.

Otherwise, let $\Phi_2 = \{F'_2, p'_2\}$. Then, the dynamic type data of B_2 is included in the data zone of the structure field F'_2 (by Lemma 19 (p. 27)), which is itself in the field data zone of B'_2 .

- If $P_1 \neq P_2$, then, by Lemma 21 (p. 30), the two field data zones of B'_1 and B'_2 are disjoint, which concludes.

Otherwise, $P_1 = P_2$ and $B'_1 = B'_2$.

- If $\Phi_1 = \emptyset$, then F and F'_2 are distinct fields (because F is scalar whereas F'_2 is a structure field) of the same subobject, so by Hypothesis (C9, p. 28) their data are disjoint.

Otherwise, let $\Phi_1 = \{F'_1, p'_1\}$ with $\text{plength}(p'_1) = \text{plength}(p_1) - 1$ to allow induction.

- If $F'_1 \neq F'_2$, then they are distinct fields of the same subobject, so their data are disjoint.
- Otherwise, we may use the induction hypothesis.

□

Disjointness of two dynamic type data The non-virtual data of two different non-virtual bases of C are not necessarily disjoint, because of those classes that share their dynamic type data with their primary bases. As this phenomenon can propagate, it is necessary to precisely determine which subobjects will share their dynamic type data.

Definition 18. *A non-virtual path l from some class B to some class A is primary if, and only if, at least one of the following conditions holds:*

- *either $B = A$ and $l = B :: \text{nil}$*
- *or B has a primary base B' and $l = B :: B' :: l'$ for some l' , such that $B' :: l'$ is a primary path from B' to A .*

Lemma 29. *For any non-virtual, non-primary path l from some class B to some class A , there exists a unique path denoted $\text{red}(l)$ and called the reduced path from l , such that there exists l' the longest primary path such that $l = \text{red}(l) @_{\text{Repeated}} l'$. By convention, if l is primary, then $\text{red}(l) \stackrel{\text{def.}}{=} B :: \text{nil}$.*

Lemma 30. *If l' is a primary path, then $\text{red}(l @_{\text{Repeated}} l') = \text{red}(l)$.*

Lemma 31. *For any non-virtual path l , $\text{nvsoff}(l) = \text{nvsoff}(\text{red}(l))$. In particular, any non-virtual path l to a dynamic class shares its dynamic type data with its reduced path $\text{red}(l)$.*

Proof. It suffices to show that for any non-virtual paths l and l' , if l' is primary, then $\text{nvsoff}(l @_{\text{Repeated}} l') = \text{nvsoff}(l)$. By induction on l' . \square

Conversely:

Lemma 32. *If two non-primary non-virtual paths l_1, l_2 from some class B to some dynamic classes A_1, A_2 such that $\text{red}(l_1) \neq \text{red}(l_2)$, then their dynamic type data are disjoint.*

Proof. Pose $l'_i = \text{red}(l_i)$ for each i . Then, as l_1 and l_2 are not primary, their reduced paths l'_1, l'_2 are not trivial. By symmetry, we may assume $\text{length}(l'_1) \leq \text{length}(l'_2)$. Then, there are two cases:

- If $l'_2 = l'_1 @_{\text{Repeated}} l$, then, as $l'_2 \neq l'_1$, l is not trivial. So, $l'_1 = \lambda + A'_1 :: A_1 :: \text{nil}$ and there are two cases:
 - either $l = A_1 :: A_2 :: \text{nil}$, so that A_2 is a non-primary direct non-virtual base of A_1 , which concludes
 - either $l = A_1 :: \lambda' :: A' :: A_2 :: \text{nil}$ for some non-virtual base A' of A_1 such that A_2 is a non-primary direct non-virtual base of A' . Then, A_2 starts at offset at least dtsize in A' , which is enough to conclude.

\square

Corollary 33. *If two non-virtual paths l_1, l_2 from some class B to some dynamic classes A_1, A_2 such that $\text{red}(l_1) \neq \text{red}(l_2)$, then their dynamic type data are disjoint.*

Proof. • If both paths were primary, then they would have the same reduced path, which is absurd.

- If (for instance) l_1 is primary and l_2 is non-primary, then A_2 starts at offset at least dtsize in B , which concludes.
- Otherwise, Lemma 32 (p. 37) concludes.

\square

Corollary 34. *If two paths $(h_1, l_1), (h_2, l_2)$ from some class C to some dynamic classes A_1, A_2 such that $\text{red}(l_1) \neq \text{red}(l_2)$, then their dynamic type data are disjoint.*

Theorem 4. *If $p_1 = (\mathcal{P}_1, j_1, P_1)$, $p_2 = (\mathcal{P}_2, j_2, P_2)$ are two generalized subobjects from $C[n]$ to some dynamic classes A_1, A_2 such that $(\mathcal{P}_1, j_1, \text{red}(P_1)) \neq (\mathcal{P}_2, j_2, \text{red}(P_2))$, then their dynamic type data are disjoint.*

Proof. As there is a condition on P_1, P_2 , we cannot use the alternate representation for pointers. So a direct reasoning is necessary.

- If $(\mathcal{P}_1, j_1) = (\mathcal{P}_2, j_2)$, then Corollary 34 (p. 38) concludes.
- Otherwise, if $\mathcal{P}_1 = \mathcal{P}_2$ and $j_1 \neq j_2$, then there are two different cells of the same array, so their data are disjoint.

Otherwise, assume $\mathcal{P}_1 \neq \mathcal{P}_2$. By symmetry, we may assume $\text{length}(\mathcal{P}_1) \leq \text{length}(\mathcal{P}_2)$. Then, we have $\mathcal{P}_2 = (j'_2, P'_2, F'_2) :: \mathcal{P}'_2$. Reason by structural induction on \mathcal{P}_1 . There are several cases:

- If $\mathcal{P}_1 = \text{nil}$, then the dynamic type data targeted by p_2 is included in the data of F'_2 which is included in the field data of the subobject p_1 . Then, Theorem 3 (p. 36) concludes.
- Otherwise, $\mathcal{P}_1 = (j'_1, P'_1, F'_1) :: \mathcal{P}'_1$. If $(j'_1, P'_1) \neq (j'_2, P'_2)$, then our generalized subobjects p_1 and p_2 are located in disjoint field zones, which concludes. Otherwise, if $F'_1 \neq F'_2$, then the data of those two fields are disjoint, which concludes. Otherwise, we may use the induction hypothesis.

□

Dynamic type data alignment

Hypothesis 25. *We assume that there exists an alignment $\text{dtdalign} > 0$ for dynamic type data.*

Then, the following theorem:

Theorem 5. *If p is a generalized subobject of static type A from a structure array of C , such that A is dynamic, then the access to the dynamic type data of A is correctly aligned:*

$$\begin{aligned} &(\text{dtdalign} \mid \text{off}_C(p)) \\ &(\text{dtdalign} \mid \text{align}_C) \end{aligned}$$

is a direct consequence of Corollary 16 (p. 24) requiring the following hypothesis:

Hypothesis 26. *If A is a dynamic class, then:*

$$(\text{dtdalign} \mid \text{nvalign}_A) \tag{C16}$$

3.3.6 Identity of subobjects

In this section, we want to show that if $p_1 \neq p_2$ are two distinct generalized subobjects of the same static type A from the same structure array of C , then they are located at distinct offsets $\text{off}_C(p_1) \neq \text{off}_C(p_2)$ within the structure array.

The proof of this theorem is very different depending on whether A is empty or not.

Non-empty base offsets If A is not empty, then (C22, p. 35) ensures that $\text{nvsize}_A \neq 0$.

However, recall that the non-virtual data of two distinct subobjects of the same static type are not necessarily disjoint, because of their dynamic type data.

Lemma 35. *If l_1, l_2 are distinct non-virtual paths from some class B to some non-empty class A , then $\text{nvsoff}(l_1) \neq \text{nvsoff}(l_2)$.*

Proof. By symmetry, we may assume $\text{length}(l_1) \leq \text{length}(l_2)$. Then, there are two cases:

- Case $l_2 = l_1 @_{\text{Repeated}} l$ is absurd. Indeed, as l_1 and l_2 have the same static type A , we would have l trivial, so $l_2 = l_1$.
- So, l_1 and l_2 point to subobjects of some classes $B'_1 \neq B'_2$ that are distinct direct non-virtual bases of some class B' . As the non-virtual data of B'_1, B'_2 are disjoint, and the non-virtual data of A is non-null, arithmetics conclude.

□

Corollary 36. *If P_1, P_2 are distinct paths from some class B to some non-empty class A , then $\text{soff}_B(P_1) \neq \text{soff}_B(P_2)$.*

Theorem 6. *If p_1, p_2 are distinct generalized subobjects from some class C to some non-empty class A , then $\text{off}_C(p_1) \neq \text{off}_C(p_2)$.*

Proof. For each i , let $p_i \propto (j_i, P_i, \Phi_i)$. By symmetry, we may assume $\text{plength}(l_1) \leq \text{plength}(l_2)$. Reason by induction on $\text{plength}(p_1)$. Then, there are several cases:

- if $j_1 \neq j_2$, then we have two distinct cells of the same array, so their data are disjoint.
- Otherwise, if $j_1 = j_2$ and $\Phi_2 = \{(F_2, p'_2)\}$, then there are two cases:
 - if $\Phi_1 = \{(F_1, p'_1)\}$, then there are three cases:
 - * if $P_1 = P_2$, then either $F_1 = F_2$, in which case we may use the induction hypothesis, or $F_1 \neq F_2$ so those are two distinct fields of the same subobject, so their data are disjoint.
 - * otherwise, the designated subobjects are located in disjoint field zones.
 - Otherwise, $\Phi_1 = \emptyset$. Then, there are several cases:
 - * Case $P_2 = P_1 @_{\text{Repeated}} P$ is absurd (the hierarchy is well-founded).

- * Case $P_1 = P_2 @_{\text{Repeated}}(B_2 :: B' :: Q)$: in that case, whereas p_2 is located in the field data zone of B_2 , p_1 is located in the data zone of B' which is a direct non-virtual base of B_2 : those two data zones are disjoint
- * Otherwise, $P_1 = L @_{\text{Repeated}} B' :: B'_1 :: Q_1$ and $P_2 = L @_{\text{Repeated}} B' :: B'_2 :: Q_2$ with B'_1 and B'_2 two distinct direct non-virtual bases of B' , so their data are disjoint.

- Otherwise, $j_1 = j_2$, $\Phi_2 = \emptyset$ and $\Phi_1 = \emptyset$ (as $\text{plength}(p_1) \leq \text{plength}(p_2)$), so Corollary 36 (p. 40) concludes.

□

Empty base offsets Unfortunately, this kind of reasoning cannot be done if A is empty. In this case, we have to explicitly enumerate the offsets to all possible subobjects of an empty static type.

Definition 19. *The sets $nveffs_C$ of the non-virtual empty base offsets of C and $effs_C$ of empty base offsets of C are defined as follows:*

$$\begin{aligned}
nveffs_C &\stackrel{\text{def.}}{=} \begin{cases} \{(C, 0)\} & \text{if } C \text{ is empty} \\ \emptyset & \text{otherwise} \end{cases} \\
&\cup \bigcup_{B \in \mathcal{NV}(C)} \text{dnvoff}_C(B) + nveffs_B \\
&\cup \bigcup_{(f, B, n) \in \text{St}\mathcal{F}(C)} \bigcup_{i \in [0..n]} \text{foff}_C(f, B, n) + i \cdot \text{size}_B + \text{effs}_B \\
\\
effs_C &\stackrel{\text{def.}}{=} \bigcup_{B \in \text{mathcal{V}}(C)} nveffs_B
\end{aligned}$$

In our Coq development, those two sets are defined as mutually inductive predicates. Only at the level of the algorithms, we show that we can construct for any class C two sets such that (A, o) is in either set if and only if (A, o) verifies the corresponding predicate, assuming that such sets exist for any class $B \prec C$. Indeed, we shall see that those sets are computed at the same time as class layout.

Lemma 37. *$(A, o) \in nveffs_C$ if, and only if, A is empty and there is a class B and a non-virtual path p from C to B such that at least one of the following conditions holds:*

- either $B = A$ and $o = \text{nvsoff}(p)$

- or there is a structure array field f of some type $B'[m']$ defined in class B and an array path \mathcal{P} from $B'[m']$ to some $A'[n']$, a nonnegative integer $k < n'$ and an inheritance path Q from A' to A such that:

$$o = \text{nvsoff}(p) + \text{aoff}_{B'}(\mathcal{P}) + k \cdot \text{size}_{A'} + \text{soff}_{A'}(Q)$$

$(A, o) \in \text{eboffs}_C$ if, and only if, A is empty and there is a class B and an inheritance path P from C to B such that at least one of the following conditions holds:

- either $B = A$ and $o = \text{soff}_C(P)$
- or there is a structure array field f of some type $B'[m']$ defined in class B and an array path \mathcal{P} from $B'[m']$ to some $A'[n']$, a nonnegative integer $k < n'$ and an inheritance path Q from A' to A such that:

$$o = \text{soff}_C(P) + \text{aoff}_{B'}(\mathcal{P}) + k \cdot \text{size}_{A'} + \text{soff}_{A'}(Q)$$

Proof. \Rightarrow : by mutual induction on the definitions of **eboffs** and **nveboffs**.

\Leftarrow : cases $B = A$ by induction on p for non-virtual inheritance, by case analysis on h where $P = (h, p)$ for virtual inheritance. Other cases by induction on \mathcal{P} . \square

Theorem 7. *Let A be an empty class. Two different pointers of static type A , but from the same initial object, point to a different memory location, if and only if, in the construction of the sets **eboffs** and **nveboffs**, the unions are disjoint, i.e. for all classes C , the following conditions hold:*

- $\forall B_1, B_2 \in \mathcal{DNV}(C) : B_1 \neq B_2 \Rightarrow$

$$\text{dnvboff}_C(B_1) + \text{nveboffs}_{B_1} \perp \text{dnvboff}_C(B_2) + \text{nveboffs}_{B_2} \quad (\text{C23})$$

- $\forall B_1 \in \mathcal{DNV}(C), \forall (f, B_2, n) \in \text{StF}(C), \forall i \in [0 \dots (n-1)] :$

$$\text{dnvboff}_C(B_1) + \text{nveboffs}_{B_1} \perp \text{foff}_C(f, B_2, n) + i \cdot \text{size}_{B_2} + \text{eboffs}_{B_2} \quad (\text{C24})$$

- $\forall (f_1, B_1, n_1), (f_2, B_2, n_2) \in \text{StF}(C), (f_1, B_1, n_1) \neq (f_2, B_2, n_2) \Rightarrow \forall i_1 \in [0 \dots (n_1-1)], \forall i_2 \in [0 \dots (n_2-1)] :$

$$\text{foff}_C(f_1, B_1, n_1) + i_1 \cdot \text{size}_{B_1} + \text{eboffs}_{B_1} \perp \text{foff}_C(f_2, B_2, n_2) + i_2 \cdot \text{size}_{B_2} + \text{eboffs}_{B_2} \quad (\text{C25})$$

- $\forall B_1, B_2 \in \tilde{\mathcal{V}}(C) : B_1 \neq B_2 \Rightarrow$

$$\text{vboff}_C(B_1) + \text{nveboffs}_{B_1} \perp \text{vboff}_C(B_2) + \text{nveboffs}_{B_2} \quad (\text{C26})$$

Proof. Mostly by definition of `nveboffs` and `eboffs`, except for two empty bases of distinct cells of a structure array. In that case, we have to use the fact that two different cells of the same array are totally disjoint (not only their data). Then, we need a further hypothesis:

Hypothesis 27. *The non-virtual size of any class B is positive.*

$$nsize_C > 0 \tag{C5}$$

□

The reason why (C5, p. 43) is required already for non-virtual sizes (not only for total sizes) may seem not obvious, but the following simple example makes it clear:

```

struct A          {};
struct B0: A      {};
struct B1: A      {};
struct C: B0, B1  {};

C c [2];
A * a0 = (A*) (B1*) &c [0];
A * a1 = (A*) (B0*) &c [1];

```

If $nsize_A = 0$, $nsize_{B_1} = 0$, and $nsize_{B_2} = 0$ were allowed, then we would not be able to distinguish pointers a_0 and a_1 above.

In practice, following the behaviour prescribed by Itanium C++ ABI [1] to compute offsets for a given class C , assuming that layout has already been computed for all bases of C and all classes that are types of some structure fields of C , components of C (base or field) are laid out one after another, and for each attempt to lay out a given component, it is checked against all components previously laid out.

Considering intervals instead of `eboffs` and `nveboffs` is not enough, even in practice. Indeed, consider the following hierarchy:

```

struct A          {};
struct B          {};
struct C: A       {};
struct D: B       {};
struct E1: A, B    {};
struct E2: A, B    {};
struct F: E1, C, E2 {};
struct G: F, D     {};

```

Then, the layout of G given by GNU C++ is such that:

$$\begin{aligned}
\text{dnvboff}_C(A) &= 0 \\
\text{dnvboff}_D(B) &= 0 \\
\text{dnvboff}_{E_1}(A) &= 0 & \text{dnvboff}_{E_1}(B) &= 0 \\
\text{dnvboff}_{E_2}(A) &= 0 & \text{dnvboff}_{E_2}(B) &= 0 \\
\text{dnvboff}_F(E_1) &= 0 & \text{dnvboff}_F(C) &= 1 & \text{dnvboff}_F(E_2) &= 2 \\
\text{dnvboff}_G(F) &= 0
\end{aligned}$$

If the type conflict test (whose existence, but not the way it should work, is prescribed by Itanium C++ ABI [1]) relies on interval checking, then base F of G would define interval $[0 \dots 2]$ for base B , so that base D of G would be laid out at offset 3. However, it would be consistent to lay out D at offset 1 within G , which GNU C++ actually does. This justifies why the exact sets of offsets have to be checked.

4 Implementation of realistic layout algorithms (CPP.v, CommonVendorABI.v, CCCPP.v)

We chose not to directly impose a layout, as Compcert does with C structures, but to set some constraints for layout algorithms, allowing them to make some optimizations that our constraints do not necessarily catch, such as field or base reordering for optimizing data alignment.

We implemented two algorithms: one based on the Itanium C++ ABI [1], and another algorithm optimized for empty base tail padding.

For both algorithms, we took a more precise definition of dynamic classes (CPP.v):

Definition 20. *A class C is dynamic if, and only if, at least one of the following conditions holds:*

- *C has a virtual method*
- *C has a virtual base*
- *C has a dynamic non-virtual direct base*

Under this definition, it is consistent to reserve no additional space for non-dynamic classes, thanks to the following limitations of the C++ Standard:

- static base-to-derived cast cannot go through virtual inheritance
- dynamic cast cannot be used from a class without virtual methods (a fortiori from a non-dynamic class)

4.1 An algorithm based on the Itanium C++ ABI

The principle of this algorithm based on the common vendor C++ ABI for Itanium is that all bases and fields are laid out within a class such that they are mutually *totally* disjoint, not only their data. This actually incurs stronger constraints than the ones stated in the previous section. So, for two such disjoint components (bases or fields), any two offsets to generalized subobjects of some class type A (empty or not) from those components are automatically distinct.

But such a layout would be too naive for empty bases, as it would consume some space for them. For this reason, an important exception is prescribed for empty bases, which may be laid out at offset 0. In that case, explicit checks are necessary to ensure that a subobject reachable from an empty base will not conflict with another subobject of the same type reachable from another base or field.

However, to algorithmically limit the need for such explicit checks, this algorithm considers classes with fields to be non-empty. More formally:

Definition 21. *A class C is empty if, and only if, all the following conditions hold:*

- *C has no virtual methods*
- *C has no direct virtual base*
- *C has no fields*
- *all direct non-virtual bases of C are empty*

In particular, a class having a virtual base is not empty. More generally:

Lemma 38. *Dynamic classes are not empty.*

Then, each class C is laid out through the following way, assuming that all classes B such that $B \prec C$ have been already laid out:

1. If C has a dynamic direct non-virtual base B , then choose B as the primary base of C , and lay out B within C at offset 0.
2. Otherwise, if C is dynamic, then reserve some space for dynamic type data
3. For each direct non-virtual base B of C that is not the primary base of C , lay it out within C as follows:
 - If B is empty, try to lay it out at offset 0, unless there is a type conflict for empty bases.

- Otherwise, try to lay out B at the current value of nvdsize_C so far. If there is a type conflict for empty bases, then increase by the non-virtual alignment nalign_B of the base B , knowing that beyond the current value for nvdsize_C , there will be no conflict (because all offsets of bases reachable from C so far are less than nvdsize_C).
 - If B is not empty, update nvdsize_C to include the *whole* non-virtual size of the base (not only its non-virtual data size).
4. Then, for each field, lay it out at increasing offsets, making them *wholly* disjoint (not only their data).
 5. Then, lay out all virtual bases of C the same way as for non-virtual bases (i.e. trying offset 0 for empty virtual bases, and making them wholly disjoint – not only their data).

During the layout of the components of C , it is also wise to construct the sets nveboffs_C and eboffs_C of generalized subobjects of an empty class type that are reachable from C .

However, the overall shape of the algorithm as described above and in the Itanium C++ ABI does not exactly state how type conflicts for empty bases should be resolved. In fact, there are only two possible conflicts:

- when trying to lay out an empty base B at offset 0, it is necessary to check that no empty base reachable through B conflict with any empty base reachable from all other non-virtual bases that have been already laid out.
- when trying to lay out a non-empty base B , or a field f , it is only necessary to check that there is no conflict with any empty base reachable from *non-virtual empty bases* laid out so far. Indeed, this algorithm lays out all non-empty components in such way that two distinct non-empty components are *totally* disjoint (not only their data).

More formally:

The main algorithm The general shape of the algorithm is as follows. Starting with the choice of the primary base and/or the allocation of dynamic type data, it then lays out the remaining components of the class, first the non-virtual bases (other than the primary base), then the fields, and finally the virtual bases.

$$\begin{aligned} \text{nvdsize}_C &\leftarrow 0, \\ \text{pbase}_C, \text{nveboffs}_C, \text{eboffs}'_C &\leftarrow \emptyset, \\ \text{dnvboff}_C, \text{foff}_C, \text{vboff}_C &: \emptyset \rightarrow \mathbb{Z}, \\ \text{nvdsize}_C, \text{nalign}_C &\leftarrow 1 \end{aligned}$$

```

if  $C$  is dynamic then
  if  $\exists B$  dynamic direct non-virtual base of  $C$  then
    pbase $_C \leftarrow \{B\}$ 
    dnvboff $_C(B) \leftarrow 0$ 
    nveboffs $_C \leftarrow$  nveboffs $_B$ 
    nvdsiz $_C \leftarrow$  nvsize $_B$ 
    nvsize $_C \leftarrow$  nvsize $_B$ 
    nvalign $_C \leftarrow$  nvalign $_B$ 
  else {there is no dynamic direct non-virtual base of  $C$ }
    nvdsiz $_C \leftarrow$  dtsize
    nvsize $_C \leftarrow$  dtsize
    nvalign $_C \leftarrow$  dtalign
  end if
end if
lay out direct non-virtual bases  $B$  of  $C$  such that pbase $_C \neq \{B\}$ 
fboundary $_C \leftarrow$  nvdsiz $_C$ 
lay out fields of  $C$ 
eboffs $_C \leftarrow$  nveboffs $_C$ 
dsiz $_C \leftarrow$  nvdsiz $_C$ 
siz $_C \leftarrow$  nvsize $_C$ 
align $_C \leftarrow$  nvalign $_C$ 
lay out virtual bases of  $C$ 
vboff $_C(C) \leftarrow 0$ 
siz $_C \leftarrow \min\{o \mid (\text{align}_C \mid o) \wedge \text{siz}_C \leq o\}$ 

```

Direct non-virtual bases

```

for all  $B$  direct non-virtual base of  $C$  such that pbase $_C \neq \{B\}$  do
  if  $B$  is empty and nveboffs $_B \perp$  nveboffs $_C$  then
    dnvboff $_C(B) \leftarrow 0$ 
  else
    dnvboff $_C(B) \leftarrow \min\{o \mid (\text{nvalign}_B \mid o) \wedge \text{nvdsiz}_C \leq o\}$ 
    while dnvboff $_C(B) < \text{nvsize}_C \wedge \neg \text{dnvboff}_C(B) + \text{nveboffs}_B \perp \underline{\text{eboffs}'_C}$ 
    do
      dnvboff $_C(B) \leftarrow$  dnvboff $_C(B) + \text{nvalign}_B$ 
    end while
  end if
  nveboffs $_C \leftarrow$  nveboffs $_C \cup$  dnvboff $_C(B) + \text{nveboffs}_B$ 
  if  $B$  is empty then
    eboffs' $_C \leftarrow$  eboffs' $_C \cup$  dnvboff $_C(B) + \text{nveboffs}_B$ 
  else
    nvdsiz $_C \leftarrow$  dnvboff $_C(B) + \underline{\text{nvsize}_B}$ 
  end if
  nvsize $_C \leftarrow \max(\text{nvsize}_C, \text{dnvboff}_C(B) + \text{nvsize}_B)$ 

```

```

    nalignC ← lcm(nalignC, nalignB)
  end for

```

Line 14 imposes every non-empty direct non-virtual base B to have its whole non-virtual part (not only its data) included in the non-virtual *data* of C . Thanks to this requirement, conflict checks for empty bases are limited to those offsets reachable only through empty direct bases (line 6). Indeed, in the “else” case of lines 5–8, B is laid out wholly (not only by its data) disjointly from all other non-empty non-virtual bases laid out so far. So, in particular, offsets to empty bases are disjoint.

Fields

```

for all  $f$  field of  $C$  do
  foffC( $f$ ) ← min{ $o$  | (falign( $f$ ) |  $o$ ) ∧ nvdsizC ≤  $o$ }
  if  $f$  is a structure field ( $fid, B, n$ ) then
    while ¬ ∪0 ≤ j < n foffC( $f$ ) +  $j$  · sizeB + eboffsB ⊥ eboffs'C do
      foffC( $f$ ) ← foffC( $f$ ) + alignB
    end while
    nveboffC ← nveboffC ∪ ∪0 ≤ j < n foffC( $f$ ) +  $j$  · sizeB + eboffsB
  end if
  nvdsizC ← foffC( $f$ ) + fsizC( $f$ )
  nvsizeC ← max(nvsizeC, foffC( $f$ ) + fsizC( $f$ ))
  nalignC ← lcm(nalignC, falign( $f$ ))
end for

```

Virtual bases

```

for all  $B$  virtual base of  $C$  do
  if  $B$  is empty and nveboffsB ⊥ eboffsC then
    vboffC( $B$ ) ← 0
  else
    vboffC( $B$ ) ← min{ $o$  | (alignB |  $o$ ) ∧ dsizC ≤  $o$ }
    while vboffC( $B$ ) < sizeC ∧ ¬vboffC( $B$ ) + nveboffsB ⊥ eboffs'C do
      vboffC( $B$ ) ← vboffC( $B$ ) + nalignB
    end while
  end if
  eboffsC ← eboffsC ∪ vboffC( $B$ ) + nveboffsB
  if  $B$  is empty then
    eboffs'C ← eboffs'C ∪ vboffC( $B$ ) + nveboffsB
  else
    dsizC ← vboffC( $B$ ) + nvsizeB
  end if
  sizeC ← max(sizeC, vboffC( $B$ ) + sizeB)
end for

```


$\text{align}_C \leftarrow \text{lcm}(\text{align}_C, \text{nvalign}_B)$
end for

Lemma 39 (Strong invariants). *This algorithm ensures that for any class C :*

$$\begin{aligned}
& \text{nvdsiz}_C \leq \text{nvsize}_C \\
& [\text{foff}_C(F_1), \text{foff}_C(F_1) + \underline{\text{fsize}}(F_1)) \\
& \perp [\text{foff}_C(F_2), \text{foff}_C(F_2) + \underline{\text{fsize}}(F_2))
\end{aligned}$$

Indeed, for this algorithm, non-empty bases and fields are laid out by ensuring that their whole sizes (not only their data) are disjoint.

Theorem 8. *This algorithm satisfies the soundness conditions C1 to C26.*

4.2 An optimized algorithm

Definition 22. *A class C is empty if, and only if, all the following conditions hold:*

- C has no virtual methods
- C has no direct virtual base
- C has no scalar fields
- if C has a structure array field of some type B , then B is empty
- all direct non-virtual bases of C are empty

Given the definition of dynamic classes taken for our algorithms, this definition for empty bases is the smallest possible satisfying the constraints given at Hypothesis 2.

Lemma 40. *Dynamic classes are not empty.*

In particular, a class having an empty virtual base is not empty.

The main algorithm The general shape of the algorithm is similar to the previous one, except when choosing a primary base: in that case, after the choice of the primary base but before the layout of the other direct non-virtual bases, the non-virtual data size is updated to the non-virtual *data* size of the primary base, not its whole non-virtual size.

$\text{nvdsiz}_C \leftarrow 0,$
 $\text{pbas}_C, \text{nveboffs}_C \leftarrow \emptyset,$
 $\text{dnvboff}_C, \text{foff}_C : \emptyset \rightarrow \mathbb{Z},$
 $\text{vboff}_C : C \mapsto 0,$
 $\text{nvsize}_C, \text{nvalign}_C \leftarrow 1$
if C is dynamic then

```

if  $\exists B$  dynamic direct non-virtual base of  $C$  then
  pbase $_C$   $\leftarrow$   $\{B\}$ 
  dnvboff $_C(B)$   $\leftarrow$  0
  nveboffs $_C$   $\leftarrow$  nveboffs $_B$ 
  nvdsiz $_C$   $\leftarrow$  nvdsiz $_B$ 
  nvsize $_C$   $\leftarrow$  nvsize $_B$ 
  nvalign $_C$   $\leftarrow$  nvalign $_B$ 
else {there is no dynamic direct non-virtual base of  $C$ }
  nvdsiz $_C$   $\leftarrow$  dtsize
  nvsize $_C$   $\leftarrow$  dtsize
  nvalign $_C$   $\leftarrow$  dtalign
end if
end if
lay out direct non-virtual bases  $B$  of  $C$  such that pbase $_C \neq \{B\}$ 
fboundary $_C$   $\leftarrow$  nvdsiz $_C$ 
lay out fields of  $C$ 
eboffs $_C$   $\leftarrow$  nveboffs $_C$ 
dsiz $_C$   $\leftarrow$  nvdsiz $_C$ 
siz $_C$   $\leftarrow$  nvsize $_C$ 
align $_C$   $\leftarrow$  nvalign $_C$ 
lay out virtual bases of  $C$ 
size $_C$   $\leftarrow$   $\min\{o \mid (\text{align}_C \mid o) \wedge \text{size}_C \leq o \wedge \text{dsiz}_C \leq o\}$ 

```

Direct non-virtual bases The algorithm is

```

for all  $B$  direct non-virtual base of  $C$  such that pbase $_C \neq \{B\}$  do
  if  $B$  is empty then
    dnvboff $_C(B)$   $\leftarrow$  0
  else
    dnvboff $_C(B)$   $\leftarrow$   $\min\{o \mid (\text{nvalign}_B \mid o) \wedge \text{nvdsiz}_C \leq o\}$ 
  end if
  while dnvboff $_C(B) < \text{nvsize}_C \wedge \neg \text{dnvboff}_C(B) + \text{nveboffs}_B \perp \text{eboffs}_C$  do
    dnvboff $_C(B)$   $\leftarrow$  dnvboff $_C(B) + \text{nvalign}_B$ 
  end while
  nveboffs $_C$   $\leftarrow$  nveboffs $_C \cup \text{dnvboff}_C(B) + \text{nveboffs}_B$ 
  if  $B$  is not empty then
    nvdsiz $_C$   $\leftarrow$  dnvboff $_C(B) + \text{nvdsiz}_B$ 
  end if
  nvsize $_C$   $\leftarrow$   $\max(\text{nvsize}_C, \text{dnvboff}_C(B) + \text{nvsize}_B)$ 
  nvalign $_C$   $\leftarrow$   $\text{lcm}(\text{nvalign}_C, \text{nvalign}_B)$ 
end for

```

Fields

```

for all  $f$  field of  $C$  do
   $\text{foff}_C(f) \leftarrow \min\{o \mid (\text{falign}(f) \mid o) \wedge \text{nvsize}_C \leq o\}$ 
  if  $f$  is a structure field  $(fid, B, n)$  then
    if  $B$  is empty then
       $\text{foff}_C(f) \leftarrow 0$ 
    else
       $\text{foff}_C(f) \leftarrow \min\{o \mid (\text{align}_B \mid o) \wedge \text{nvsize}_C \leq o\}$ 
    end if
    while  $\neg \bigcup_{0 \leq j < n} \text{foff}_C(f) + j \cdot \text{size}_B + \text{eboffs}_B \perp \text{eboffs}_C$  do
       $\text{foff}_C(f) \leftarrow \text{foff}_C(f) + \text{align}_B$ 
    end while
     $\text{nveboff}_C \leftarrow \text{nveboff}_C \cup \bigcup_{0 \leq j < n} \text{foff}_C(f) + j \cdot \text{size}_B + \text{eboffs}_B$ 
    if  $B$  is not empty then
       $\text{nvsize}_C \leftarrow \text{foff}_C(f) + \text{fsize}(f)$ 
    end if
    else  $\{f$  is a scalar field $\}$ 
       $\text{nvsize}_C \leftarrow \text{foff}_C(f) + \text{fsize}(f)$ 
    end if
     $\text{nvsize}_C \leftarrow \max(\text{nvsize}_C, \text{foff}_C(f) + \text{fsize}(f))$ 
     $\text{nvalign}_C \leftarrow \text{lcm}(\text{nvalign}_C, \text{falign}(f))$ 
  end for

```

Virtual bases

```

for all  $B$  virtual base of  $C$  do
  if  $B$  is empty then
     $\text{vboff}_C(B) \leftarrow 0$ 
  else
     $\text{vboff}_C(B) \leftarrow \min\{o \mid (\text{align}_B \mid o) \wedge \text{dsize}_C \leq o\}$ 
  end if
  while  $\text{vboff}_C(B) < \text{size}_C \wedge \neg \text{vboff}_C(B) + \text{nveboffs}_B \perp \text{eboffs}_C$  do
     $\text{vboff}_C(B) \leftarrow \text{vboff}_C(B) + \text{nvalign}_B$ 
  end while
   $\text{eboffs}_C \leftarrow \text{eboffs}_C \cup \text{vboff}_C(B) + \text{nveboffs}_B$ 
  if  $B$  is not empty then
     $\text{dsize}_C \leftarrow \text{vboff}_C(B) + \text{nvsize}_B$ 
  end if
   $\text{size}_C \leftarrow \max(\text{size}_C, \text{vboff}_C(B) + \text{size}_B)$ 
   $\text{align}_C \leftarrow \text{lcm}(\text{align}_C, \text{nvalign}_B)$ 
end for

```

References

- [1] Itanium C++ ABI. <http://www.codesourcery.com/public/cxx-abi/>, March 2001.
- [2] Sandrine Blazy and Xavier Leroy. Mechanized semantics for the Clight subset of the C language. *Journal of Automated Reasoning*, 43(3):263–288, 2009.
- [3] Juan Chen. A typed intermediate language for compiling multiple inheritance. *SIGPLAN Not.*, 42(1):25–30, 2007.
- [4] The C++ Standard Committee. ISO/IEC 14882:2003 – The C++ Programming Language.
- [5] Nathan C. Myers. The "Empty Member" C++ Optimization. *C++ Issue, Dr. Dobb's Journal*, 1997.
- [6] Tahina Ramananandro, Gabriel Dos Reis, and Xavier Leroy. Formal verification of object layout for c++ multiple inheritance. Submitted to POPL 2011.
- [7] Jonathan G. Rossie, Jr., and Daniel P. Friedman. An algebraic semantics of subobjects. In *Proceedings of the Tenth Annual Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA '95)*, pages 187–199. ACM Press, 1996.
- [8] Bjarne Stroustrup. *A history of C++: 1979–1991*, pages 699–769. ACM, New York, NY, USA, 1996.
- [9] Daniel Wasserrab, Tobias Nipkow, Gregor Snelting, and Frank Tip. An operational semantics and type safety proof for multiple inheritance in C++. In *OOPSLA '06: Proceedings of the 21st annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications*, pages 345–362, New York, NY, USA, 2006. ACM.