

# Frama-C Value Analysis

## Séminaire CAP'TRONIC

Virgile Prevosto  
virgile.prevosto@cea.fr

June 18<sup>th</sup>, 2015



long ra  
for 0 =>  
C1) if (m  
tmp2 =  
of the

tmp2[0] = 1 << (nbl - 1) else if (tmp1[0]) >= 1 << (nbl - 1) tmp2[0] = 1 << (nbl - 1) + abs(tmp2[0] - tmp1[0]); /\* Then the second part looks like the first one: \*/  
tmp1[0] = 0; k = 0; k << 8; k << 1; tmp1[0] += mc2[0][k] \* tmp2[k]; /\* The [i][j] coefficient of the matrix product MC2\*TMP2, that is: \*MC2\*(TMP1) = MC2\*(MC1\*M1) = MC2\*M1 \*MC1  
i = 1; tmp1[0] >= 1; /\* Final rounding: tmp2[0] is now represented on 9 bits: \*if (tmp1[0] < -256) m2[0] = -256; else if (tmp1[0] > 255) m2[0] = 255; else m2[0] = tmp1[0];

## Introduction

## Abstract domains

- Arithmetic
- Memory

## Methodology

- Basic commands
- Parameters

```
(long n)
for (i = 0; i < n; i++)
  tmp2[i] = 1;
// ...
// ...
```

```
tmp2[i] = (i < (n-1) ? tmp1[i] : 1); // Then the second part looks like the first one:
tmp1[i+k] = 0; k = 5; k++) tmp1[i+k] += mc2[i][k] * tmp2[k]; // The [i][k] coefficient of the matrix product MC2*TMP2, that is: *MC2*(TMP1) = MC2*(MC1*M1) = MC2*M1*MC1
i = 1; tmp1[0] >= 1; // Final rounding: tmp2[0] is now represented on 9 bits: if (tmp1[0] < -256) m2[0] = -256; else if (tmp1[0] > 255) m2[0] = 255; else m2[0] = tmp1[0];
```



## Credits

- ▶ Pascal Cuoq
- ▶ Boris Yakobowski
- ▶ Matthieu Lemerre
- ▶ A few other developers...

## More information

- ▶ <http://frama-c.com/download/frama-c-value-analysis.pdf>



## Find the domains of the variables of a program

- ▶ based on **abstract interpretation**
- ▶ **alarms** on operations that **may** be invalid
- ▶ alarms on the specifications that may be invalid
- ▶ **Correct**: if no alarm is raised, no runtime error can occur

```

long n;
for (i = 0; i < n; i++)
  tmp2 =
  // ...

```


```

tmp2[i][k] = (i < 2 * (n-1) ? tmp2[i][k] : (i < 2 * (n-1) ? tmp2[i][k] : tmp2[i][k] + tmp2[i][k]) / 2; // Then the second part takes the first part
tmp1[i][k] = 0; k = 5; k = tmp1[i][k] + mc2[i][k] * tmp2[i][k]; // The [i][k] coefficient of the matrix product MC2 * TMP2, that is, *MC2*(TMP1) = MC2*(MC1 * M1) = MC2 * M1 * MC1
i = i + tmp1[i][k] >= 1; // Final rounding: tmp2[i][k] is now represented on 9 bits. *if (tmp1[i][k] < -256) tmp2[i][k] = -256; else if (tmp1[i][k] > 255) tmp2[i][k] = 255; else tmp2[i][k] = tmp1[i][k];

```



# Abstract Interpretation in a Nutshell



- ▶ Replace all possible concrete execution ...
- ▶ ... by one abstract execution
- ▶ Analysis is guaranteed to terminate
- ▶ Over-approximation and false alarms
- ▶ Trade-off between precision and computation time

```

(long n)
for (i = 0; i < n; i++)
    tmp2 = ...
    
```


```

tmp2[0] = 1;
for (k = 0; k < n; k++)
    tmp1[k] = 0;
for (i = 0; i < n; i++)
    for (k = 0; k < n; k++)
        tmp1[i][k] = m2[0][k] * tmp2[k];
    
```

The [i][j] coefficient of the matrix product MC2\*TMP2, that is, \*MC2\*[i][TMP2] = MC2\*[i][M1] \* MC1[i][j]. Final rounding: tmp2[0] is now represented on 9 bits. \*M1[tmp1][0] < 256? m2[0][0] = 256; else m2[0][0] = tmp1[0][0];



# Abstract Interpretation in a Nutshell



- ▶ Replace all possible concrete execution ...
- ▶ ... by one abstract execution
- ▶ Analysis is guaranteed to terminate
- ▶ Over-approximation and false alarms
- ▶ Trade-off between precision and computation time


(long ra  
t for 0 =>  
C1) if (m  
tmp2  
se of the

tmp2[0] = 1; if ((nbl - 1) && !tmp1[0]) >= 1; if ((nbl - 1) && tmp2[0]) = tmp1[0]; /\* Then the second part looks like the first one. \*/  
tmp1[0][0] = 0; k = 0; k++ tmp1[0][k] += mc2[0][k] \* tmp2[0][k]; /\* The [i,j] coefficient of the matrix product MC2\*TMP2, that is, \*MC2\*(TMP1) = MC2\*(MC1\*M1) = MC2\*M1 \*MC1  
i = 1; tmp1[0][i] >= 1; /\* Final rounding: tmp2[0][i] is now represented on 9 bits. \*if (tmp1[0][i] < -256) m2[0][i] = -256; else if (tmp1[0][i] > 255) m2[0][i] = 255; else tmp2[0][i] = tmp1[0][i];




# Abstract Interpretation in a Nutshell

- ▶ Replace all possible concrete execution ...
- ▶ ... by one abstract execution
- ▶ Analysis is guaranteed to terminate
- ▶ Over-approximation and false alarms
- ▶ Trade-off between precision and computation time



# Abstract Interpretation in a Nutshell

- ▶ Replace all possible concrete execution ...
- ▶ ... by one abstract execution
- ▶ Analysis is guaranteed to terminate
- ▶ **Over-approximation and false alarms**
- ▶ Trade-off between precision and computation time



```

long ra
for (i = 0; i < n; i++)
    tmp2[i] = tmp1[i] * 2;

```

```


tmp2[0] = 0;
for (k = 0; k < n; k++)
    tmp1[k] = m2[k] * tmp2[k];
}
The [i, j] coefficient of the matrix product MC2 * TMP2, that is, *MC2*(TMP2) = MC2 * (M1 * MC1) = MC2 * M1 * MC1.
Final rounding: tmp2[0] is now represented on 9 bits. If (tmp1[0] < 256) m2[0] = 256; else if (tmp1[0] < 255) m2[0] = 255; else m2[0] = 254;

```





# Abstract Interpretation in a Nutshell



- ▶ Replace all possible concrete execution ...
- ▶ ... by one abstract execution
- ▶ Analysis is guaranteed to terminate
- ▶ Over-approximation and false alarms
- ▶ Trade-off between precision and computation time

```

long n;
for (i = 0; i < n; i++)
  tmp2[i] = 0;

```

```

tmp2[0] = 0;
for (k = 0; k < n; k++)
  tmp1[k] = 0;
for (j = 0; j < n; j++)
  tmp2[j] = 0;
for (i = 0; i < n; i++)
  for (k = 0; k < n; k++)
    tmp1[k] = m2[0][k] * tmp2[k];

```

The [i][j] coefficient of the matrix product  $MC2 * TMP2$ , that is,  $*MC2[i](TMP2) = MC2[i](MC1 * M1) = MC2 * M1[i](MC1 * M1) = MC2 * M1[i] * M1$ .



## Some specificities

- ▶ Precise handling of **pointers**
- ▶ Several representation for dynamic allocation (precision vs. time)
- ▶ time and memory efficient (as much as achievable)
- ▶ Precise enough
  - ▶ for proving absence of runtime errors on some critical code
  - ▶ to serve as a back-end for other semantical analyzes through its API



# Integer and Floating Point Arithmetic

## Corresponding Abstract Domain

small set of integers (by default, cardinal  $\leq 8$ )

- ⊕ integer interval  $\times$  modulo information
- ⊕ finite floating-point interval

## Examples

- ▶  $\{0; 40; \}$  = 0 or 40
- ▶  $[0..40]$  = an integer between 0 and 40 (inclusive)
- ▶  $[-..-]$  = any integer (within the bound of the corresponding integral type)
- ▶  $[3..39], 3\%4$  = 3, 7, 11, 15, 19, 23, 27, 31, 35 or 39
- ▶  $[0.25..3.125]$  = floating-point between 0.25 and 3.125 (inclusive)

▶ next



```

int x, y, t, m; double d;
extern char z; char z1;

void f(int c) {
    if (c) x = 40;
    for (int i = 0; i<=40; i++) {
        Frama_C_show_each_loop_1(i);
        if (c == i) y = i; }
    z1 = z;
    t = z;
    m=3;
    for (int i = 3; i<=40; i+=4) {
        if (c == i) m = i; }
    if (c) { d = 0.25; } else { d = 3.125; }
}
  
```



```
frama-c -val -main f integer.c
```

```
[value] Called Frama_C_show_each_loop_1({0; 1})
[value] Called Frama_C_show_each_loop_1({0; 1; 2})
[value] Called Frama_C_show_each_loop_1([0..16])
[value] Called Frama_C_show_each_loop_1([0..40])
[value] ===== VALUES COMPUTED =====
  x IN {0; 40}
  y IN [0..40]
  z1 IN [--..--]
  t IN [-128..127]
  m IN [3..39], 3%4
  d IN [0.25 .. 3.125]
```



## Base Address

Global variable

- ⊕ Formal parameter of main function
- ⊕ literal string constant
- ⊕ NULL
- ⊕ ...

## Addresses

- ▶ Base address + Offset (integer)
- ▶ Each base has a maximal valid offset
- ▶ Abstract Values are sets of addresses

▶ next



# Examples of Addresses

## Precise Base

- ▶  $\{\{\&p + \{4; 8\}\}\}$  = address of  $p$  shifted from 4 or 8 octets
- ▶  $\{\{\&"foobar";\}\}$  = Address of literal string "foobar" (shifted from 0)
- ▶  $\{\{\&NULL + \{1024;\}\}\}$  = Absolute location 1024

## Imprecision

- ▶ garbled mix of  $\&\{x_1; \dots; x_n\}$  = unknown address built upon arithmetic operations over integers and addresses  $x_1; \dots; x_n$ .
- ▶ **ANYTHING** = top of the lattice. Should not occur in practice



```
int* x,*z, *t; const char* y; int p[3];
const char* string = "foobar";
```

```
void f(int c) {
  if (c) { x = &p[1]; }
  else { x = &p[2]; }
  y = string;
  z = (int*)1024;
  t = (int*) ((int)x | 4096);
}
```





```

[value] ===== VALUES COMPUTED =====
[value] Values at end of function f:
  x IN {{ &p{[1], [2]} }}
  y IN {{ "foobar" }}
  z IN {1024}
  t IN
  {{ garbled mix of &p{
    (origin: Arithmetic
     examples/value/address.c:16)} }}
  
```



# Write to an Address

## Abstract Domain

written address = valid left value

address

- × initialized?
- × *not dangling pointer?*

## Example

```

{
  int x, y;
  if (e) x = 2;
L: if (e) y = x + 1;
}
  
```

- ▶ At *L*, we know that *x* equals 2 iff it has been initialized
- ▶ Depending on the complexity of *e*, we know that *y* equals 3 if *x* equals 2



```

int X, Y, *p;
void f(int c) {
    int x, y;
    if (c<=0) x = 2;
    L: if (c<=0) y = x + 1; else y = 4;
    X = x;
    Y = y;
    p = c ? &X : &x;
}

```

```

int main(int c) {
    f(c);
    if (Y==4) *p = 3;
    return 0;
}

```




```
examples/value/address_written.c:8:
[kernel] warning:
    accessing uninitialized left-value:
    assert \initialized(&x);
examples/value/address_written.c:16:
[kernel] warning:
    accessing left-value that
    contains escaping addresses:
    assert !\dangling(&p);
[value] Values at end of function main:
X IN {2; 3} or UNINITIALIZED
Y IN {3; 4}
p IN {{ &X }} or ESCAPINGADDR
__retres IN {0}
```



# Concrete Memory

- ▶ Seen as big array of bits
- ▶ read/write a value  $v$  at address  $i =$  read/write  $v$  at index  $i$  over  $n$  bits.
- ▶  $n$  depends upon the type of  $v$
- ▶ potential **overlap**
- ▶ **example:**  $x.a$  extends over  $n$  bits,  $x.b$  over  $m$  bits. Writing at  $x.a$  over  $p$  bytes with  $n < p < m$  will partially erase  $x.b$



# Abstract Memory State

base address  $\rightarrow$  offsetmap

## Example

$$S \mapsto \{ [0..31] \mapsto \{ \&x \mapsto 0 \text{ (initialized, not dangling)} \}$$

$$[32..47] \mapsto \{ \text{NULL} \mapsto 12 \text{ (initialized, not dangling)} \} \}$$

$$x \mapsto \{ [0..31] \mapsto \{ \text{NULL} \mapsto \{ 3; 24 \} \text{ (initialized, not dangling)} \} \}$$

## Displaying Values

- ▶ **Expected type** is used to display values

## Exemple

$$S.\text{mypointer} \in \{ \{ \&x \} \}$$

$$.\text{myshort} \in 12$$

$$x \in \{ 3; 24 \}$$

▶ next



# Abstract Memory

```

struct S {
    int* mypointer;
    short myshort;
} S;
char T[sizeof(struct S)];
int x;
void main (int c) {
    if(c) { x = 3; } else { x = 24; }
    S.mypointer = &x;
    S.myshort = 12;
    *(int*)T = &x;
    *(short *)((int*)T + 1) = 12;
}
  
```



```
[value] ===== VALUES COMPUTED =====
[value] Values at end of function main:
  S.mypointer IN {{ &x }}
  .myshort IN {12}
  .[bits 48 to 63] IN {0}
  T[bits 0 to 31] IN {{ (? *)&x }}
  [bits 32 to 47] IN {12}
  [6..7] IN {0}
  x IN {3; 24}
```



long n...  
for 0 <= k <= 5  
C1[i][k] =  
tmp2...  
of the

tmp2[0][i] = 0; k = 5; k++) tmp1[0][i] += mc2[0][k] \* tmp2[k][i]; /\* The [i][j] coefficient of the matrix product MC2\*TMP2, that is, \*MC2\*(TMP1) = MC2\*(MC1\*M1) = MC2\*M1\*MC1...  
l = 1; tmp1[0][i] >>= 1; \*/ Final rounding: tmp2[0][i] is now represented on 9 bits. \*if (tmp1[0][i] < -255) tmp2[0][i] = -255; else if (tmp1[0][i] > 255) tmp2[0][i] = 255; else tmp2[0][i] = tmp1[0][i];



# Overlapping

```
int c, x;
```

```
char t[6];
```

```
void main(void) {
    t[0] = c ? 1 : 2;
    *(int*)(t+1) = c ? 3 : 4;
    *(t+3) = 5;
    x = *(int*)(t+1);
}
```



```

[value] Values at end of function main:
  x[bits 0 to 15]# IN {4}%32, bits 0 to 15
    [bits 16 to 23] IN {5}
    [bits 24 to 31]# IN {4}%32, bits 24 to 31
  t[0] IN {2}
    [bits 8 to 23]# IN {4}%32, bits 0 to 15
    [3] IN {5}
    [4]# IN {4}%32, bits 24 to 31
    [5] IN {0}
  
```



# Main options

- ▶ `-main`: specifies the entry point of the analysis (default: `main` function)
- ▶ `-lib-entry`: Library mode: globals are not assumed to be 0-initialized
- ▶ `-val`: launch value, starting at the specified entry point

`abs.c`

*// returns the absolute value of x*

```
int abs ( int x ) {
    if ( x >=0 )
        return x ;
}
```

- ▶ `frama-c -main abs -val -lib-entry abs.c`



## Is abstract interpretation an automated plug-in?

- ▶ yes...
- ▶ and no!
- ▶ **must be driven** carefully to give meaningful results
- ▶ requires some expertise and some time



# Help Value Analysis to Understand the Code

- ▶ Pay attention to **missing code** (external library) or code that is **not understood** (asm)
  - ▶ write **C code** (stub), that can be understood by Value and approximates the missing part well enough with respect to the desired property
  - ▶ give an **ACSL specification**
- ▶ Give an appropriate **context**
  - ▶ Write an appropriate **entry point** to initialize global variables and formal parameters
  - ▶ Sometime possible to use **dedicated options** (`-context-*`)

▶ next



# Code Sample

```

void getchar(char* p);

int main(char* string) {
    char c;
    int idx = 0;
    getchar(&c);
    while(string[idx]) {
        if (string[idx] == c)
            return idx;
        idx++;
    }
    return -1;
}

```



# Value result without specification

```

examples/value/stub.c:7:
[kernel] warning: out of bounds read.
    assert \valid_read(string+idx);
examples/value/stub.c:8: [kernel] warning:
    accessing uninitialized left-value:
    assert \initialized(&c);
examples/value/stub.c:8:
[kernel] warning: out of bounds read.
    assert \valid_read(string+idx);
[value] Values at end of function main:
    c IN [--..--] or UNINITIALIZED
    idx IN {0; 1}
    __retres IN {-1; 0; 1}
  
```



# Value result with specification

[value] Values at end of function main:

```
c_0 IN [--..--]
idx IN [0..255]
__retres IN [-1..254]
```

[value] Values at end of function real\_main:

```
s[0..254] IN [--..--]
[255] IN {0}
```

long ra  
for 0 <=  
C1) if (m  
tmp2 =  
of the

tmp2[0] = 1 <= (n-1) else if (tmp1[0]) >= 1 <= (n-1) - 1; else tmp2[0] = tmp1[0]; 7. Then the second part takes like the first part. 8. tmp1[0] = 0; k = 5; k = k + 1; tmp1[0] = m2[0][k] \* tmp2[0][k]; 9. The [i][j] coefficient of the matrix product MC2 \* TMP2, that is, \*MC2\*(TMP1) = MC2\*(MC1 \* M1) = MC2 \* M1 \* MC1. 10. i = 1; tmp1[0] >= 1; 11. Final rounding: tmp2[0] is now represented on 9 bits. \*if (tmp1[0] < -256) tmp2[0] = -256; else if (tmp1[0] > 255) tmp2[0] = 255; else tmp2[0] = tmp1[0];





# Enhance precision

## Loops

- ▶ option `-ulevel`: syntactic loop unrolling
- ▶ option `-slevel`: allows Value to explore  $n$  separated paths before joining them
- ▶ option `-wlevel`: number of loop steps before performing widening (default is 3, use with caution)

## Driving Value through Annotations

- ▶ **ACSL assertions** can be used to restrict propagated domains
- ▶ but only if Value can interpret it

```
/*@ assert x % 2 == 0; */
```

```
// potentially useful
```

```
/*@ assert \exists integer y; x == 2 * y; */
```

```
// useless
```

- ▶ Case analysis using **disjunctions**



# Code Sample

```

int x, y;

void main (int c) {
    if (c) { x = 10; } else { x = 33; }
    if (!c) { x++; } else { x--; }

    if (c<=0) { y = 42; } else { y = 36; }
    if (c>0) { y++; } else { y--; }
}

```

long n;  
for (i = 0; i < n; i++)  
 tmp2[i] = 0;  
 ...

tmp2[0] = 1; for (k = 1; k < n; k++) tmp1[k] = mc2[0][k] \* tmp2[k];  
The [i][j] coefficient of the matrix product MC2\*TMP2, that is, \*MC2\*(TMP1) = MC2\*(MC1\*M1) = MC2\*M1\*MC1  
= 1 \* tmp1[0][i] >= 1; Final rounding: tmp2[0][i] is now represented on 9 bits: \*if (tmp1[0][i] < -256) tmp2[0][i] = -256; else if (tmp1[0][i] > 255) tmp2[0][i] = 255; else tmp2[0][i] = tmp1[0][i];



## without level

```
x IN {9; 11; 32; 34}
y IN {35; 37; 41; 43}
```

## with level, no assertion

```
x IN {9; 11; 34}
y IN {37; 41}
```

## with level and assertion

```
/*@ assert c<=0 || c > 0; */
```

[value] Assertion got status valid.

```
x IN {9; 34}
y IN {37; 41}
```



# A more complete example

- ▶ Demo on PolarSSL 1.3.7 code
- ▶ Actually detected by Trust in Soft (<http://trust-in-soft.com/>), a spin-off of CEA
- ▶ 12 bugs reported and fixed since v1.1.8 according to PolarSSL Changelog

