

CEA Contribution & Future Work

Julien Signoles

CEA LIST

Software Safety Labs

Hi-Lite Final Meeting

May 2013, the 29th



list

HI-LITE

Simplifying the use of formal methods



long ra
for B =>
C1) if (a
tmp2 =
of the

tmp2[0] = 1 << (nbl - 1) else if (tmp1[0]) >= 1 << (nbl - 1) tmp2[0] = tmp1[0]; 7. Then the second part takes the first one: tmp1[0] = 0; k = 0; k <= 5; k++) tmp1[0][k] = mc2[0][k] * tmp2[0][k]; 7. 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) tmp2[0][i] = -256; else if (tmp1[0][i] > 255) tmp2[0][i] = 255; else tmp2[0][i] = tmp1[0][i];

1. what's the status before Hi-Lite (Apr. 2010)
2. what's done during Hi-Lite
 - ▶ E-ACSL language
 - ▶ E-ACSL plug-in
 - ▶ improving Frama-C
3. what's planned to do next

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

tmp2[i][j] = 0; if (i <= (nbl - 1) && j <= (nbl - 1)) tmp2[i][j] = (i <= (nbl - 1) && j <= (nbl - 1)) ? tmp2[i][j] + tmp1[i][j] : 0; Then the second part takes like the first one.
tmp1[i][j] = 0; k = 0; k <= 5; k++) tmp1[i][j] += mc2[i][k] * tmp2[k][j]; /* The [i][j] coefficient of the matrix product MC2*TMP2, that is, *MC2*(TMP1) = MC2*(MC1*M1) = MC2*M1 *MC1
i = 1; tmp1[i][j] >>= 2; */ Final rounding: tmp2[i][j] is now represented on 9 bits. *if (tmp1[i][j] < -256) m2[i][j] = -256; else if (tmp1[i][j] > 255) m2[i][j] = 255; else tmp1[i][j] =



- ▶ **Frama-C** Boron
- ▶ only designed for **static analyzers**
- ▶ mostly **no way to combine** them to verify program properties
- ▶ **ACSL** specification language
- ▶ **PathCrawler** test generation tool

long no
 for 0 =>
 C1) if (m
 tmp2 =
 re of the

tmp2[0] = 1; for (k=0; k<N; k++) tmp1[k] += mc2[0][k] * tmp2[k]; /* The [j] coefficient of the matrix product MC2*TMP2, that is, *MC2*(TMP1) = MC2*(MC1*M1) = MC2*M1*MC1
 l = 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];



- ▶ **Frama-C** Fluorine: Boron + 4
- ▶ now designed both for **static and dynamic analyzers**
- ▶ **combining analyses** is effective
- ▶ **ACSL + E-ACSL** specification languages
- ▶ **PathCrawler** as a Frama-C Plug-in

What was done during these 3 years in Hi-Lite?



long no
for 0 =>
C1) if (m
tmp2 =
of the
tmp2[0][i] = 0; k = 0; k++) tmp1[0][i] += m2[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
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];

Testing tools:

- ▶ require a precondition specifying valid inputs
- ▶ require an oracle to decide whether a test is correct

Abstract interpreters:

- ▶ require a precondition and assertions to be precise

Program proving tools:

- ▶ require a formal specification
- ▶ based on pre/post-conditions

Combining them requires a **common specification language**



- ▶ **E-ACSL: Executable ANSI/ISO C Specification Language**
 - ▶ builds a bridge between static and dynamic analysis tools
 - ▶ based on pre-existing **ACSL language** used by Frama-C

- ▶ **E-ACSL plug-in converts E-ACSL specifications into C code**
 - ▶ Frama-C plug-in
 - ▶ runtime assertion checking
 - ▶ helpful for debugging specification
 - ▶ may easily be used by any analysis tool for C

long n;
for (i = 0; i < n; i++)
C[i] = 0;
tmp2 = 0;
for (k = 0; k < n; k++)
tmp1[k] = 0;

tmp2[0] = 1; for (i = 1; i < n; i++) tmp2[i] = tmp2[i-1] + tmp1[i-1];
for (k = 0; k < n; k++) tmp1[k] = mc2[0][k] * tmp2[k];
Final rounding: tmp2[0] is now represented on 3 bits: if (tmp1[0] < 255) tmp2[0] = 255; else if (tmp1[0] > 255) tmp2[0] = 255; else tmp2[0] = tmp1[0];



Benefits:

- ▶ being executable allows to be **understandable by dynamic tools** (testing tools, monitors)
- ▶ being based on ACSL allows to be **supported by existing Frama-C analyzers**
- ▶ being translatable into C allows to be **supported by other analysis tools for C**

Differences with ACSL:

- ▶ few restrictions
- ▶ one extension: **iterators** over recursive datastructures
- ▶ compatible semantics changes



- ▶ **quantifications** must be guarded

```
\forall x1, ..., xn;
  a1 <= x1 <= b1 && ... && an <= xn <= bn
  ==> p
```

```
\exists x1, ..., xn;
  a1 <= x1 <= b1 && ... && an <= xn <= bn
  && p
```

- ▶ **sets** must be finite
- ▶ **loop invariants** are simply equivalent to 2 assertions
- ▶ no way to express **termination** properties
- ▶ backwards C **labels** only



Iterators over C recursive datastructures

```

// type of binary trees
struct btree {
    int val;
    struct btree *left, *right;
};

// declare an iterator over a binary tree
/*@ iterator access(_, struct btree *t):
    @ nexts t->left, t->right;
    @ guards \valid(t->left), \valid(t->right); */

// is_even(t) is valid iff all values in the binary tree t are even
/*@ predicate is_even(struct btree *t) =
    @ \forallall struct btree *tt;
    @ access(tt, t) ==> tt->val % 2 == 0; */

```



- ▶ **mathematical integers** to preserve ACSL semantics
- ▶ many advantages compared to bounded integers
 - ▶ **automatic theorem provers** work much better with such integers than with bounded integers arithmetics
 - ▶ specify **without implementation details in mind**
 - ▶ still **possible to use bounded integers** when required
 - ▶ much easier to **specify overflows**

long n
for 0 <=
c1; if (n
tmp2 =
of the

tmp2[0] = 1; for (int i = 1; i < n; i++) tmp2[i] = tmp2[i-1] * 2; // Then the second part takes the first part
tmp1[0] = 0; for (int i = 1; i < n; i++) tmp1[i] = mc2[0][i] * tmp2[i]; // The [i][j] coefficient of the matrix product MC2*TMP2, that is, *MC2*(TMP1) = MC2*(MC1*M1) = MC2*M1*MC1
if (tmp1[0] >= 1) { // Final rounding: tmp2[0] is now represented on 3 bits: if (tmp1[0] <= 256) tmp2[0] = 256; else if (tmp1[0] > 255) tmp2[0] = 255; else tmp2[0] = tmp1[0];



- ▶ ACSL logic is total and $1/0$ is logically significant
 - ▶ help the user to write simple specification like $u/v == 2$
 - ▶ $1/0$ is defined but not executable

- ▶ E-ACSL logic is 3-valued
 - ▶ the semantics of $1/0$ is “undefined”
 - ▶ lazy operators $\&\&$, $||$, $_{-?}_{-}:$, $==>$
 - ▶ correspond to Chalin’s Runtime Assertion Checking semantics

 - ▶ consistent with ACSL: valid (resp. invalid) E-ACSL predicates remain valid (resp. invalid) in ACSL



- ▶ convert E-ACSL annotations into C code
- ▶ implemented as a Frama-C plug-in

```
int div(int x, int y) {
  /*@ assert y-1 != 0; */
  return x / (y-1);
}
```

E-ACSL

```
int div(int x, int y) {
  /*@ assert y-1 != 0; */
  e_acsl_assert(y-1 != 0);
  return x / (y-1);
}
```



- ▶ convert E-ACSL annotations into C code
- ▶ implemented as a Frama-C plug-in

<pre>int div(int x, int y) { /*@ assert y-1 != 0; */ return x / (y-1); }</pre>	<p>E-ACSL →</p>	<pre>int div(int x, int y) { /*@ assert y-1 != 0; */ e_acsl_assert(y-1 != 0); return x / (y-1); }</pre>
--	-----------------	---

- ▶ the general translation is more complex than it may look
 - ▶ `\result` requires to introduce extra-variables
 - ▶ `\at(x,L)` requires to introduce code at L
 - ▶ ...



- ▶ use **GMP library** for mathematical integers

```

/*@ assert y-1 == 0; */
mpz_t e_acsl_1, e_acsl_2, e_acsl_3, e_acsl_4;
int e_acsl_5;
mpz_init_set_si(e_acsl_1, y); // e_acsl_1 = y
mpz_init_set_si(e_acsl_2, 1); // e_acsl_2 = 1
mpz_init(e_acsl_3);
mpz_sub(e_acsl_3, e_acsl_1, e_acsl_2); // e_acsl_3 = y-1
mpz_init_set_si(e_acsl_4, 0); // e_acsl_4 = 0
e_acsl_5 = mpz_cmp(e_acsl_3, e_acsl_4); // (y-1) == 0
e_acsl_assert(e_acsl_5 == 0); // runtime check
mpz_clear(e_acsl_1); mpz_clear(e_acsl_2); // deallocate
mpz_clear(e_acsl_3); mpz_clear(e_acsl_4);

```



- ▶ use **GMP library** for mathematical integers

```

/*@ assert y-1 == 0; */
mpz_t e_acsl_1, e_acsl_2, e_acsl_3, e_acsl_4;
int e_acsl_5;
mpz_init_set_si(e_acsl_1, y); // e_acsl_1 = y
mpz_init_set_si(e_acsl_2, 1); // e_acsl_2 = 1
mpz_init(e_acsl_3);
mpz_sub(e_acsl_3, e_acsl_1, e_acsl_2); // e_acsl_3 = y-1
mpz_init_set_si(e_acsl_4, 0); // e_acsl_4 = 0
e_acsl_5 = mpz_cmp(e_acsl_3, e_acsl_4); // (y-1) == 0
e_acsl_assert(e_acsl_5 == 0); // runtime check
mpz_clear(e_acsl_1); mpz_clear(e_acsl_2); // deallocate
mpz_clear(e_acsl_3); mpz_clear(e_acsl_4);

```

- ▶ design a **type system** to detect when GMP is really required
- ▶ infer a correct **interval** for any term, as small as possible
- ▶ almost **no GMP in practice** :-)



must prevent introducing RTE when translating annotations

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  return u/v;
}
```



must prevent introducing RTE when translating annotations

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  return u/v;
}
```

E-ACSL

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```

long ra
for 0 <=
C1; if (u
tmp2 =
of the

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



must prevent introducing RTE when translating annotations

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  return u/v;
}
```

E-ACSL

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```

RTE plug-in

```
int foo(int u, int v) {
  /*@ assert v != 0; */
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```



must prevent introducing RTE when translating annotations

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  return u/v;
}
```

E-ACSL

```
int foo(int u, int v) {
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```

RTE plug-in

```
int foo(int u, int v) {
  /*@ assert v != 0; */
  e_acsl_assert(v != 0);
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```

E-ACSL

```
int foo(int u, int v) {
  /*@ assert v != 0; */
  /*@ assert u/v == 2; */
  e_acsl_assert(u/v == 2);
  return u/v;
}
```



- ▶ memory-related constructs like `\valid` require to know the memory structure at runtime
- ▶ C library for memory observation
- ▶ used by E-ACSL Plug-in
- ▶ once again the translation is quite heavy
- ▶ **backward dataflow analysis** to instrument the code only when required



- ▶ J. Signoles.
E-ACSL User Manual.
May 2013.
- ▶ M. Delahaye, N. Kosmatov and J. Signoles.
Common Specification Language for Static and Dynamic
Analysis of C Programs.
SAC'13. March 2013.
- ▶ N. Kosmatov, G. Petiot and J. Signoles.
Optimized Memory Monitoring for Runtime Assertion
Checking of C Programs.
Submitted article.
- ▶ N. Kosmatov and J. Signoles.
Runtime Assertion Checking with Frama-C.
Submitted tutorial.



- ▶ how to ensure the **safety of an annotated program**
- ▶ by using several **customizable analyzers**
- ▶ based on **different techniques?**
- ▶ a **“consolidation algorithm”** merges all the results coming from the different analyzers with their different configurations
- ▶ potential results are:
 - ▶ **valid**
 - ▶ **unknown**
 - ▶ **invalid**
 - ▶ **inconsistent**
 - ▶ a variety of refinement (never tried, dead annotations, ...)



► the consolidation algorithm is correct

if each analyzer is correct, then the algorithm returns “Valid” (resp. “Invalid”) for a valid (resp. invalid) property. It returns “Inconsistent” if there are both a proof of validity and invalidity.

► the consolidation algorithm is complete

if each analyzer is correct and indicates the right hypotheses, and if one analyzer does not indicate “Dont know” under recursively valid hypotheses, then the computed status is either “Valid” or “Invalid”.



- ▶ L. Correnson and J. Signoles.
Combining Analyses for C Program Verification.
 FMICS'12. Aug. 2012.
- ▶ P. Cuoq, F. Kirchner, N. Kosmatov, V. Prevosto, J. Signoles
 and B. Yakobowski.
Frama-C, A Software Analysis Perspective.
 SEFM'12. Oct. 2012.
 Selected for journal publication.
- ▶ P. Cuoq, D. Doligez and J. Signoles.
Lightweight Typed Customizable Unmarshaling.
 ML'11. Sep. 2011.



- ▶ support **missing constructs**:
 - ▶ **assigns** and loop assigns
 - ▶ **logic functions** and **predicates**
 - ▶ **loop invariants**
 - ▶ **complete** and **disjoint behaviors**
 - ▶ ...
- ▶ **temporal memory safety** (balancing of malloc/free, ...)
- ▶ memory **profiling**
- ▶ improve the instrumentation: more **optimizations**
- ▶ **proof** of E-ACSL optimized instrumentation



- ▶ E-ACSL was initially designed for runtime assertion checking
- ▶ debugging specifications
 - ▶ before proving program
 - ▶ teaching
- ▶ monitoring
 - ▶ security application
 - ▶ combining monitoring and static analysis
 - ▶ demo this afternoon
- ▶ combining test and static analysis



- ▶ 1 opened **Phd position**
 - ▶ Formalization of **E-ACSL** within **Coq**
- ▶ 1 submitted French **ANR project**
 - ▶ combining static and dynamic analyses
 - ▶ fully centered around **E-ACSL**
- ▶ 1 European **Artemis** project being submitted
 - ▶ security-oriented
 - ▶ E-ACSL for **monitoring on a simulator**
- ▶ **Sec4Safe**
 - ▶ when, where, what, who? :-)
- ▶ Other projects?



- ▶ **Tool** collaborations
- ▶ **Language** Collaborations
 - ▶ Mixed C/Ada program verification
- ▶ Tools/analysis/language Collaborations in a **certification** context
 - ▶ 1 opened **Phd position** (combining test and proof)

SRI's **Evidential Tool Bus**



long n
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]; /* Then the second part takes the first part
tmp1[0] = 0; k = 5; k--> tmp1[0] += m2[0][k] * tmp2[k]; /* The [k] coefficient of the matrix product M2*TMP2, that is, * M2*(TMP1) = M2*(M1 * M1) = M2*M1 * M1
l = 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 tmp2[0] = tmp1[0];

- ▶ **E-ACSL**: new executable specification language for C
- ▶ implemented as a **Frama-C plug-in**
- ▶ **combining analysis is now effective** within Frama-C
- ▶ 3 articles + 1 submitted, 1 short paper, 1 submitted tutorial
- ▶ **several potential applications**
- ▶ a lot of works remain (both theoretical and practical)



(long n) for 0 <= k < n
 C1) if (n <= 0) tmp2 = 0;
 of the
 tmp2[0] = 1; else if (tmp1[0] >= 0) tmp2[0] = 1; else if (tmp1[0] < 0) tmp2[0] = 0;
 Then the second part takes the first part
 tmp1[0] = 0; k = 1; while (k < n) tmp1[k] = m2[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
 l = 1; tmp1[0][l] >= 1; Final rounding: tmp2[0][l] is now represented on 9 bits: if (tmp1[0][l] < -256) tmp2[0][l] = -256; else if (tmp1[0][l] > 255) tmp2[0][l] = 255; else tmp2[0][l] = tmp1[0][l];