

Through the Steps of Programmable Refactoring of a Large Scientific Code

Michele MARTONE


Leibniz Supercomputing Centre
Garching bei München, Germany

Università degli studi di Verona
November 25, 2019



“ The GADGET simulation code¹


- ★ ☆ Cosmological large-scale structure formation (galaxies and clusters)
- 🔗 ① Highly scalable ($O(100k)$ Xeon cores on SuperMUC@LRZ)
- 👥 🧑🏫 Several teams and versions (>100 kLoC each)



¹V. Springel, “The cosmological simulation code GADGET-2”, MNRAS, vol. 364, pp. 1105–1134, 2005

Q Performance pilot study at LRZ

- ▶ L.Iapichino, V. Karakasis, F.Baruffa, N.Hammer
- ▶ spanned over one year
- ▶ focused on 1kLoC extract
- ▶ identified changes meant for whole GADGET



Speedup requires data layout change \Leftrightarrow ²

main header:

```
1 struct particle {
2     double Mass, ...
3 #if defined(BLACK_HOLES)...
4     double Hsml, ...
5     ...
6 };
```



```
1 struct particle_soa_t {
2     double *Mass, ...
3 #if defined(BLACK_HOLES)...
4     double *Hsml, ...
5     ...
6 };
```

init source file:

```
1 // Array of Structures:
2 struct particle *P;
3 // allocate one global array:
4 P = mymalloc(...
5
6
7
8
```



```
1 // Structure of Arrays
2 struct particle_soa_t P_SoA;
3 // allocate one global array
4 // for each field:
5 P_SoA.Mass = mymalloc(...
6 #if defined(BLACK_HOLES)...
7 P_SoA.Hsml = mymalloc(...
8 ...
```

favour auto-vectorization in *.c:

```
1 ...
2 // may not vectorize
3 P[i].Mass + P[i]...
```



```
1 ...
2 // vectorizes better
3 P_SoA.Mass[i] + P_SoA...
```

²F.Baruffa et al. <https://arxiv.org/abs/1612.06090>

Speedup suggests major code change steps

1. each type, $10 \approx 100$ fields, `#ifdefs`:

```
1 struct particle {
2     double Mass, ...
3     #if defined(BLACK_HOLES)...
4     double Hsml, ...
5     ...
6 };
```



```
1 struct particle_soa_t {
2     double *Mass, ...
3     #if defined(BLACK_HOLES)...
4     double *Hsml, ...
5     ...
6 };
```

2. (almost) each field an allocation:

```
1 // Array of Structures:
2 struct particle *P;
3 // allocate one global array:
4 P = mymalloc(...
5
6
7
8
```



```
1 // Structure of Arrays
2 struct particle_soa_t P_SoA;
3 // allocate one global array
4 // for each field:
5 P_SoA.Mass = mymalloc(...
6 #if defined(BLACK_HOLES)...
7 P_SoA.Hsml = mymalloc(...
8 ...
```



3. $\gg 10KLoC$ change in *.c!


```
1 ...
2 // may not vectorize
3 P[i].Mass + P[i]...
```



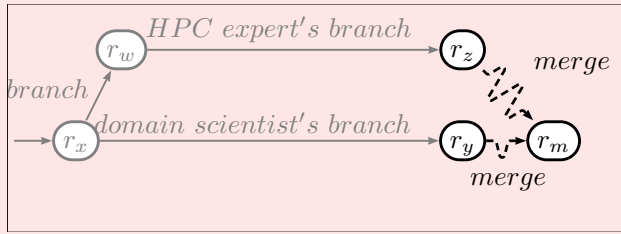
```
1 ...
2 // vectorizes better
3 P_SoA.Mass[i] + P_SoA...
```

» 10KLoC diff = 

 ~~by hand~~? 

 custom script?

 and ...all-at-once?



A C code matching and transformation engine



- ▶ A project from Inria (France)



²<https://git.kernel.org/pub/scm/linux/kernel/git/backports/backports.git/tree/patches>

³<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/scripts/coccinelle>

A C code matching and transformation engine



- ▶ A project from Inria (France)
- ▶ originally to
 -  update Linux kernel drivers²
 -  **smash bugs**
(hence the name)³



²<https://git.kernel.org/pub/scm/linux/kernel/git/backports/backports.git/tree/patches>

³<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/scripts/coccinelle>

A C code matching and transformation engine

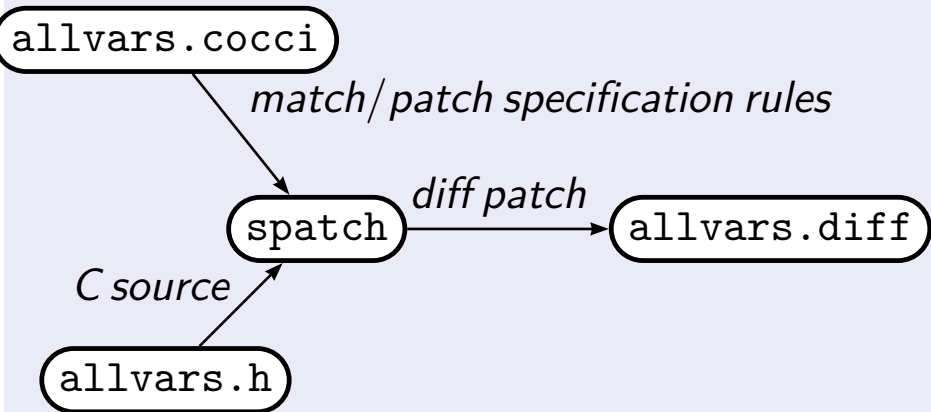
- ▶ A project from Inria (France)
- ▶ originally to
 -  update Linux kernel drivers²
 -  **smash bugs**
(hence the name)³
- ▶ seemingly **underutilized** in other contexts



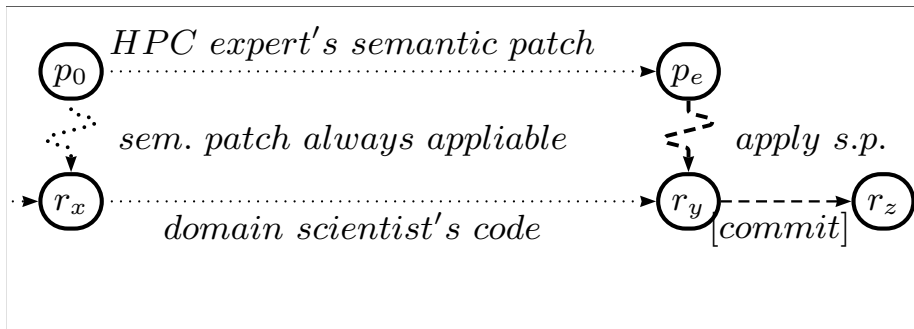
²<https://git.kernel.org/pub/scm/linux/kernel/git/backports/backports.git/tree/patches>

³<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/scripts/coccinelle>

Patch what and how ?




Transformations *encoded* in *semantic patches*



- ▶ produces usable patched (e.g. refactored) code
- ▶ no requirement to commit patched code

Step 1:

modify data
structures into SoA



Match main data structure

```
1 @prctl_str@
2 identifier id = {particle_data, sph_particle_data};
3 field list fs;
4 identifier I;
5 declaration d;
6 type ST;
7 @@
8
9 (
10  struct id { fs } *I;
11 &
12  ST          { fs } *I;
13 &
14  d
15 )
```

- ▶ pattern language to *match C*
- ▶ *metavariables* `id`, `I`, ... to match C language elements

Create new identifier

```
1 @script:python new_prtcl_str_id@
2 id << prtcl_str.id;
3 id1;
4 @@
5
6 coccinelle.id1="%s_soa_t"%(id)
```

- ▶ reuse `identifier` `id` from rule `prtcl_str`
- ▶ new string using Python

Clone main data structure

```
1 @insert_new_prtcl_str depends on prtcl_str@
2 identifier new_prtcl_str_id.id1;
3 field list prtcl_str.fs;
4 type T;
5 @@
6
7 extern int maxThreads;
8 ++struct id1 { fs };
```

- ▶ pattern language to *patch* C
- ▶ reminiscent of GNU patch

Filter out unwanted fields ▼

```
1 @match_anon_union_in_struct@
2 identifier id;
3 identifier J;
4 field list[n] fs;
5 identifier new_ptcl_str_id.id1;
6 @@
7
8 struct id1 { fs
9   union { ... } J;
10   ...
11 };
12
13 @rm_union_from_struct depends on match_anon_union_in_struct@
14 field list[match_anon_union_in_struct.n] fs;
15 identifier new_ptcl_str_id.id1;
16 field fld;
17 @@
18
19 struct id1 { fs
20 - fld
21   ...
22 };
```


Whitelisted types to pointers


```
1 @make_ptr@
2 identifier new_ptrcl_str_id.id1;
3 identifier M;
4 typedef MyDouble;
5 typedef MyFloat;
6 typedef MyLongDouble;
7 typedef MyDoublePos;
8 typedef MyBigFloat;
9 type MT = { double, float, MyDouble, MyFloat, MyLongDouble,
            MyDoublePos, MyBigFloat};
10
11
12 @@
13
14 struct id1 { ...
15 - MT M;
16 + MT *M;
17   ...
18 };
```

Remove non-pointer fields ✘

```
1 @del_non_ptr@
2 identifier new_ptrcl_str_id.id1;
3 identifier J;
4 type T;
5 type P != {T*};
6 @@
7
8 struct id1 { ...
9 - P J;
10   ...
11 };
```

► for each field `J` of type `P` which is not a pointer to another type...

Step 2: allocation functions



Populate allocation functions ↗

```
1 @per_type_soa_alloc@
2 identifier new_prtcl_str_id.id1;
3 identifier prtcl_str_mmbrrs.M;
4 type prtcl_str_mmbrrs.MT;
5 symbol P;
6 identifier insert_per_type_soa_functions.soa_alloc_fid;
7 fresh identifier si = "#ifdef "##HAVE__"##id1##"__"##M##" //";
8
9 identifier str_from_id.str;
10 @@
11 void soa_alloc_fid(...)
12 { ...
13 ++si;
14 ++      P->M = (MT*) mymalloc(str, sizeof(*(P->M)) * N);
15 ++      if(!P->M)soa_abort(/*"allocating "*/ str);
16 ++#endif
17 }
```

Step 3:

AoS → SoA



Match old structure defs⁴

```
1 @ostr@
2 identifier id = {particle_data, sph_particle_data};
3 type ST;
4 identifier P;
5 @@
6
7 (
8   struct id { ... } *P;
9 &
10  ST { ... } *P;
11 )
```

⁴Subset of rule `prtcl_str` from p. 10.

Match fields of new types⁵

```
1 @nt@
2 identifier pps.id1;
3 identifier I;
4 type T;
5 @@
6
7 struct id1 {
8     ...
9     T I;
10    ...
11 };
```

⁵Relies on rule pps, duplicate of new_prtcl_str_id, p. 11.

Mold new identifiers⁶

```
1 @script:python pid@
2 id1 << pps.id1;
3 P << ostr.P;
4 S;
5 @@
6
7 coccinelle.S="%s_soa"%(P)
```

⁶Same mechanism as `new_prtcl_str_id` at p. 11.

Patch many expressions ($\gg 10\text{KLoC}$ diff)

```
1 @soa_access@
2 identifier ostr.P;
3 identifier pid.S;
4 identifier nt.I;
5 expression E;
6 @@
7
8 - P[E].I
9 + S.I[E]
```

\forall `identifier` P previously matched in rule `ostr`

\forall `identifier` S previously matched in rule `pid`

\forall `identifier` I previously matched in rule `nt`

\forall `expression` E matching rule `soa_access`

substitute `P[E].I` with `S.I[E]`

Outcome ✓


- 🚀 GADGET can evolve further
 - ▶ its semantic patches...
 - ▶ ...can be stored 📁
 - ▶ ...and applied at anytime ▶
 - ▶ ...ease performance experiments 🧪
 - ▶ ...serve also as a Coccinelle test⁷ 🔄

Thanks to Dr. Julia Lawall (Inria) for successful collaboration!

⁷<https://github.com/coccinelle/coccinelle/commit/ad5a94>

One-day training: 24.03.2020 at LRZ

Introduction to Semantic Patching of C programs with Coccinelle



Register online

https://www.lrz.de/services/compute/courses/2020-03-24_hspc2w19/