

Adjoint Data-Flow analyses applied to checkpointing - Tradeoff between snapshots and TBR

Benjamin Dauvergne

TROPICS Project, INRIA Sophia-Antipolis

Why checkpoints?

- Instead of recording the tape of the execution, you want to reexecute some part of your code.
- To do this you need to restore the variables used by this part to the value they carried at the time of the first execution.
- **Used** here means read before written, it is a classical data flow analysis notation, like **Def**.

$$\mathbf{Use}(I_1, \dots, I_n) = \mathbf{Use}(I_1) \cup (\mathbf{Use}(I_2, \dots, I_n) \setminus \mathbf{Def}(I_1))$$

.

Usual way of doing checkpoints

- By hand :
we know the code, we know that there is something called the *state* and it is read and written between checkpoints. We create a procedure which saves it on the tape and we provide it to the AD tool.
- Automatically:
when you write a source to source AD tool you don't know what the input code is doing, so you need data flow analysis to find out those used variables and if they will be overwritten.

What should we save?

Data flow notation from a previous paper of L. Hascoet and M. Araya.

$$\begin{aligned} X &= [I_1, \dots, I_n] \text{ a sequence of instructions} \\ \text{adjoint program of } X &= \emptyset \vdash \bar{X} \text{ where} \\ \mathbf{TBR} \vdash \bar{I}; \bar{D} &= \left| \begin{array}{l} \mathbf{PUSH}(\mathbf{Def}(I) \cap (\mathbf{TBR} \cup \mathbf{Use}(I'))) \\ I \\ (\mathbf{TBR} \cup \mathbf{Use}(I')) \setminus \mathbf{Def}(I) \vdash \bar{D} \\ \mathbf{POP}(\mathbf{Def}(I) \cap (\mathbf{TBR} \cup \mathbf{Use}(I'))) \\ I' \end{array} \right. \end{aligned}$$


- I' is the adjoint code associated with a single instruction. When you differentiate you have a context: save set **TBR**.

The TBR - Snapshot trade off

Bigger TBR

$$\begin{aligned} \text{TBR} \vdash \overline{C;D} &= \text{PUSH}(\text{Def}(C) \cap \text{TBR}) \\ &\quad \text{PUSH}(\text{Def}(C) \cap \text{Use}(\overline{C})) \\ &\quad C \\ &\quad (\text{TBR} \cup \text{Use}(\overline{C})) \setminus \text{Def}(C) \vdash \overline{D} \\ &\quad \text{POP}(\text{Def}(C) \cap \text{Use}(\overline{C})) \\ &\quad \emptyset \vdash \overline{C} \\ &\quad \text{POP}(\text{Def}(C) \cap \text{TBR}) \end{aligned}$$

Bigger Snapshot

$$\begin{aligned} \text{TBR} \vdash \overline{C;D} &= \text{PUSH}(\text{Def}(C) \cap \text{TBR}) \\ &\quad \text{PUSH}(\text{Def}(C;D) \cap \text{Use}(\overline{C})) \\ &\quad C \\ &\quad \text{TBR} \setminus (\text{Def}(C) \cup \text{Snap}) \vdash \overline{D} \\ &\quad \text{POP}(\text{Def}(C;D) \cap \text{Use}(\overline{C})) \\ &\quad \emptyset \vdash \overline{C} \\ &\quad \text{POP}(\text{Def}(C) \cap \text{TBR}) \end{aligned}$$


A code where «big snapshots» are bad

Loop | $proc_1(\text{Use } state, \text{Def } A)$
| $proc_2(\text{Use } state, \text{Def } B)$
| $proc_3(\text{Use } state, \text{Def } C)$
| $proc_4(\text{Use } ABC, \text{Def } state)$

In Tapenade we checkpoint all calls so this example is interesting.

A code where «big snapshots» are bad

The forward sweep of preceding code using «big snapshots».

Loop	PUSH (<i>state</i>)
	<i>proc</i> ₁ (Use <i>state</i> , Def A)
	PUSH (<i>state</i>)
	<i>proc</i> ₂ (Use <i>state</i> , Def B)
	PUSH (<i>state</i>)
	<i>proc</i> ₃ (Use <i>state</i> , Def C)
	PUSH (A, B, C)
	<i>proc</i> ₄ (Use ABC, Def <i>state</i>)

It's not really good, each time we save *state*, we save the same values.

A code where «big snapshots» are bad

The forward sweep of preceding code using «big TBR».

Loop	PUSH (A)
	<i>proc</i> ₁ (Use <i>state</i> , Def A)
	PUSH (B)
	<i>proc</i> ₂ (Use <i>state</i> , Def B)
	PUSH (C)
	<i>proc</i> ₃ (Use <i>state</i> , Def C)
	PUSH (<i>state</i>)
	<i>proc</i> ₄ (Use <i>ABC</i> , Def <i>state</i>)

Now we are able to remove redundant **PUSH**.

A code where « big TBR » is bad

$proc_1(use = arrayA)$

a gather/scatter loop on A

- The forward sweep of preceding code using «big TBR»:

$proc_1(use = arrayA)$

a gather/scatter loop on A full of **PUSH**(A(i))

#**PUSH** > sizeof(A).

- The forward sweep of preceding code using «big snapshots»:

PUSH(A)

$proc_1(use = arrayA)$

a gather/scatter loop on A with less **PUSH**

Numerical results

On one of our test code using the « big snapshots » scheme:

Time of original function:	2.269999962300062
Time of tangent AD function:	7.000000000000000
Time of reverse AD function:	25.48999786376953
Max Stack size:	15876 blocks of 16384 bytes

with a always « big TBR » scheme :

Time of original function:	2.289999943226576
Time of tangent AD function:	7.090000152587891
Time of reverse AD function:	22.73000049591064
Max Stack size:	11815 blocks of 16384 bytes

It's a 26% gain in terms of memory and a 11% gain on cpu, without even knowing the code.

Conclusion

- It is important to look at how you compute your snapshots.
- «big TBR» is the scheme which gives the better result in general.
- If a static analysis can infer that an array is going to be completely written once or more just after, «big snapshots» seems to be appropriate.

Further work

- Find more, easily detectable code patterns, where one or the other scheme is better.
- How could flow dependant data flow informations help us ?
i.e specialization at run-time or using profiling.
- Array region analysis.
- The placement of checkpoints in big callgraphs/flowgraphs.