

Managing Performance vs. Accuracy Trade-offs with an Improved Bit-Level Precision Tuning

CMMSE 2021

Assalé Adjé Dorra Ben Khalifa Matthieu Martel
dorra.ben-khalifa@univ-perp.fr

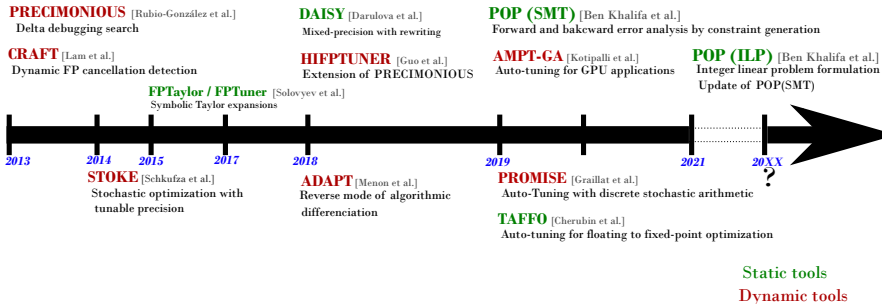
July 25, 2021

- ▶ Higher-precision data formats, e.g. *binary64* and *binary128*
 - + More accurate/robust
 - Higher execution time, increased memory / energy consumption



- ▶ **Challenge:** Best trade-offs between accuracy and performance
- ▶ **Precision Tuning:** lowering a computational kernel's precision up to a user accuracy requirement
- ▶ **Idea:** develop automated techniques to assist in precision tuning

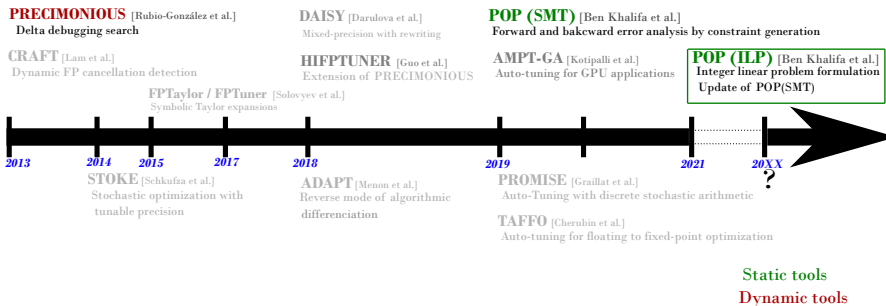
Precision Tuning: A Survey



- ▶ Try and fail strategies: change more or less randomly the data types and run the program

[S. Cherubin and G. Agosta. Tools for Reduced Precision Computation: A Survey. In ACM Computing Surveys'20]

Precision Tuning: A Survey




[D. Ben Khalifa, M. Martel, and A. Adjé. POP : A Tuning Assistant for Mixed-Precision Floating-Point Computations (FTSCS'19)]

[A. Adjé, D. Ben Khalifa, and M. Martel. Fast and Efficient Bit-Level Precision Tuning (SAS'21)]


Preliminary Notations and Definitions

ufp, nsb, ulp and computation errors



Preliminary Notations and Definitions

ufp, nsb, ulp and computation errors



- The **unit in the first place** of a real number x

$$\text{ufp}(x) = \begin{cases} \min\{i \in \mathbb{Z} : 2^{i+1} > |x|\} = \lfloor \log_2(|x|) \rfloor & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

Preliminary Notations and Definitions

ufp, nsb, ulp and computation errors



- ▶ The **unit in the first place** of a real number x

$$\text{ufp}(x) = \begin{cases} \min\{i \in \mathbb{Z} : 2^{i+1} > |x|\} = \lfloor \log_2(|x|) \rfloor & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$


- ▶ $\text{nsb}(x)$: number of significant bits of x
 - ▶ \hat{x} : approximation of x in finite precision
 - ▶ $\varepsilon(x) = |x - \hat{x}|$: the absolute error

$$\varepsilon(x) \leq 2^{\text{ufp}(x) - \text{nsb}(x) + 1}$$

[Parker, D.S.: Monte carlo arithmetic: exploiting randomness in floating-point arithmetic. Tech. Rep. CSD-970002, University of California'97]

Preliminary Notations and Definitions


ufp, nsb, ulp and computation errors



- ▶ The **unit in the last place** of x : $ulp(x) = ufp(x) - nsb(x) + 1$
- ▶ $ufp_e(x) = ufp(x) - nsb(x)$

Preliminary Notations and Definitions

ufp, nsb, ulp and computation errors




- ▶ The **unit in the last place** of x : $ulp(x) = ufp(x) - nsb(x) + 1$
- ▶ $ulp_e(x) = ufp(x) - nsb(x)$
- ▶ nsb_e : number of significant bits of the computation error on x
- ▶ $nsb_e(x)$ is computed as follows:
 - ▶ For a constant c , $nsb_e(c) = 0$ (constants assumed exact)
 - ▶ $x^\ell = c_1^{\ell_1} \odot c_2^{\ell_2}$ with $ulp_e(c_1) \geq ulp_e(c_2)$, $\odot \in \{+, -, \times, \div\}$

$$nsb_e(x) = ulp_e(c_1) - (ulp_e(c_2) - nsb_e(c_2) + 1)$$

Preliminary Notations and Definitions

ufp, nsb, ulp and computation errors




- ▶ The **unit in the last place** of x : $ulp(x) = ufp(x) - nsb(x) + 1$
- ▶ $ufp_e(x) = ufp(x) - nsb(x)$
- ▶ nsb_e : number of significant bits of the computation error on x
- ▶ $nsb_e(x)$ is computed as follows:
 - ▶ For a constant c , $nsb_e(c) = 0$ (constants assumed exact)
 - ▶ $x^\ell = c_1^{\ell_1} \odot c_2^{\ell_2}$ with $ufp_e(c_1) \geq ufp_e(c_2)$, $\odot \in \{+, -, \times, \div\}$

$$nsb_e(x) = ufp_e(c_1) - (ufp_e(c_2) - nsb_e(c_2) + 1)$$

Preliminary Notations and Definitions

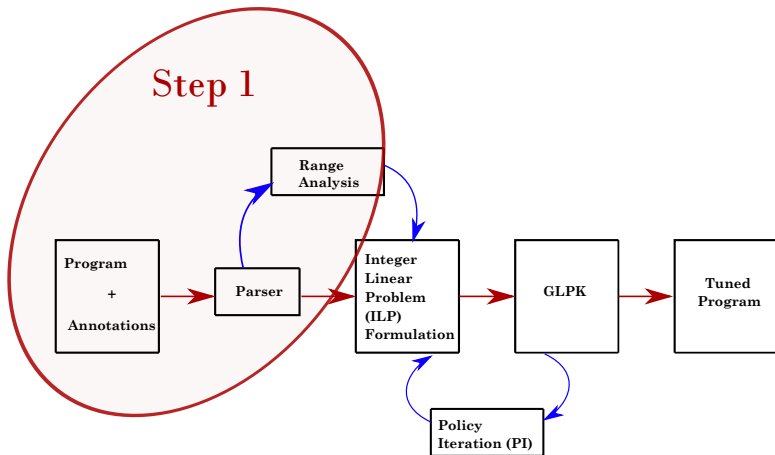
ufp, nsb, ulp and computation errors



- ▶ The **unit in the last place** of x : $ulp(x) = ufp(x) - nsb(x) + 1$
- ▶ $ufp_e(x) = ufp(x) - nsb(x)$
- ▶ nsb_e : number of significant bits of the computation error on x
- ▶ $nsb_e(x)$ is computed as follows:
 - ▶ For a constant c , $nsb_e(c) = 0$ (constants assumed exact)
 - ▶ $x^\ell = c_1^{\ell_1} \odot c_2^{\ell_2}$ with $ufp_e(c_1) \geq ufp_e(c_2)$, $\odot \in \{+, -, \times, \div\}$
$$nsb_e(x) = ufp_e(c_1) - (ufp_e(c_2) - nsb_e(c_2) + 1)$$
- ▶ $ulp_e(x) = ufp_e(x) - nsb_e(x) + 1$


POP: Precision OPTimizer

Step 1: Parsing and Range Determination Phase



Parsing and Range Determination Phase

Pendulum Example




```
1 g = 9.81; l = 0.5;
2 y1 = 0.785398;
3 y2 = 0.785398;
4 h = 0.1; t = 0.0;
5 while (t < 10.0) {
6   y1new = y1 + y2 * h ;
7   aux1 = sin(y1) ;
8   aux2 = aux1 * h * g / l ;
9   y2new = y2 - aux2;
10  t = t + h;
11  y1 = y1new;
12  y2 = y2new;
13 };
14 require_nsb(y2, 20);
```

POP Source File

Parsing and Range Determination Phase

Pendulum Example



```
1 g = 9.81; l = 0.5;
2 y1 = 0.785398;
3 y2 = 0.785398;
4 h = 0.1; t = 0.0;
5 while (t < 10.0) {
6   y1new = y1 + y2 * h ;
7   aux1 = sin(y1) ;
8   aux2 = aux1 * h * g / l ;
9   y2new = y2 - aux2 ;
10  t = t + h ;
11  y1 = y1new ;
12  y2 = y2new ;
13 };
14 require_nsb(y2, 20);
```

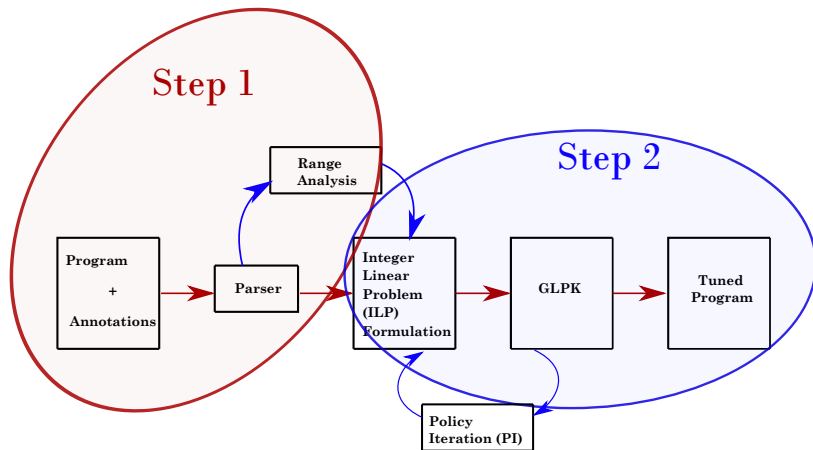
POP Source File

```
1 gℓ1 = 9.81ℓ0; lℓ3 = 0.5ℓ2;
2 y1ℓ5 = 0.785398ℓ4;
3 y2ℓ7 = 0.785398ℓ6;
4 hℓ9 = 0.1ℓ8; tℓ11 = 0.0ℓ10;
5 while (tℓ13 <ℓ15 10.0ℓ14)ℓ59 {
6   y1newℓ24 = y1ℓ17 +ℓ23 y2ℓ19 *ℓ22 hℓ21;
7   aux1ℓ28 = sin(y1ℓ26)ℓ27;
8   aux2ℓ40 = aux1ℓ30 *ℓ39 hℓ32
9     *ℓ38 gℓ34 /ℓ37 lℓ36;
10  y2newℓ46 = y2ℓ42 -ℓ45 aux2ℓ44;
11  tℓ52 = tℓ48 +ℓ51 hℓ50;
12  y1ℓ55 = y1newℓ54;
13  y2ℓ58 = y2newℓ57; };
14 require_nsb(y2, 20)ℓ61;
```

POP Label File

Constraints Generation for Bit-Level Tuning

Step 2: ILP Formulation




Constraints Generation for Bit-Level Tuning

ILP Formulation with Pessimistic Carry Bit Propagation



- ▶ Static technique relying on semantical modelling of the propagation of numerical errors
- ▶ Reasoning on ulp_e (known at constraint generation time) and nsb (unknown) of the program values
- ▶ Generate an ILP from the program source code
- ▶ Optimally solved by a classical LP solver in **polynomial time**
- ▶ Pessimistic carry bit function: $\xi = 1$



Constraints Generation for Bit-Level Tuning

Pendulum Program (ILP)



```
1 gℓ1 = 9.81ℓ0; |ℓ3 = 0.5ℓ2;
2 y1ℓ5 = 0.785398ℓ4;
3 y2ℓ7 = 0.785398ℓ6;
4 hℓ9 = 0.1ℓ8; tℓ11 = 0.0ℓ10;
5 while (tℓ13 <ℓ15 10.0ℓ14)ℓ59 {
6   y1newℓ24 = y1ℓ17 +ℓ23 y2ℓ19 *ℓ22 hℓ21;
7   aux1ℓ28 = sin(y1ℓ26)ℓ27;
8   aux2ℓ40 = aux1ℓ30 *ℓ39 hℓ32
9     *ℓ38 gℓ34 /ℓ37 |ℓ36;
10  y2newℓ46 = y2ℓ42 -ℓ45 aux2ℓ44;
11  tℓ52 = tℓ48 +ℓ51 hℓ50;
12  y1ℓ55 = y1newℓ54;
13  y2ℓ58 = y2newℓ57; };
14 require_nsb(y2,20)ℓ61;
```

POP Label File

Constraints Generation for Bit-Level Tuning

Pendulum Program (ILP)



```
1 gℓ1 = 9.81ℓ0; lℓ3 = 0.5ℓ2;
2 y1ℓ5 = 0.785398ℓ4;
3 y2ℓ7 = 0.785398ℓ6;
4 hℓ9 = 0.1ℓ8; tℓ11 = 0.0ℓ10;
5 while (tℓ13 <ℓ15 10.0ℓ14)ℓ59 {
6   y1newℓ24 = y1ℓ17 +ℓ23 y2ℓ19 *ℓ22 hℓ21;
7   aux1ℓ28 = sin(y1ℓ26)ℓ27;
8   aux2ℓ40 = aux1ℓ30 *ℓ39 hℓ32
9   *ℓ38 gℓ34 /ℓ37 lℓ36;
10  y2newℓ46 = y2ℓ42 -ℓ45 aux2ℓ44;
11  tℓ52 = tℓ48 +ℓ51 hℓ50;
12  y1ℓ55 = y1newℓ54;
13  y2ℓ58 = y2newℓ57; }
14 require_nsb(y2,20)ℓ61;
```

$$C_1 = \begin{cases} \text{nsb}(\ell_{17}) \geq \text{nsb}(\ell_{23}) + (-1) + \xi(\ell_{23})(\ell_{17}, \ell_{22}) - (-1) // \text{ADD} \\ \text{nsb}(\ell_{22}) \geq \text{nsb}(\ell_{23}) + 0 + \xi(\ell_{23})(\ell_{17}, \ell_{22}) - (1) // \text{ADD} \\ \text{nsb}(\ell_{19}) \geq \text{nsb}(\ell_{22}) + \xi(\ell_{22})(\ell_{19}, \ell_{21}) - 1 // \text{MULT} \\ \text{nsb}(\ell_{21}) \geq \text{nsb}(\ell_{22}) + \xi(\ell_{22})(\ell_{19}, \ell_{21}) - 1 // \text{MULT} \\ \text{nsb}(\ell_{23}) \geq \text{nsb}(\ell_{24}) // \text{ASSIGN} \\ \xi(\ell_{23})(\ell_{17}, \ell_{22}) \geq 1, \xi(\ell_{22})(\ell_{19}, \ell_{21}) \geq 1 // \text{CARRY BIT} \end{cases}$$

POP Label File

Constraints Generation for Bit-Level Tuning

Pendulum Program (ILP)



```
1 gℓ1 = 9.81ℓ0; lℓ3 = 0.5ℓ2;
2 y1ℓ5 = 0.785398ℓ4;
3 y2ℓ7 = 0.785398ℓ6;
4 hℓ9 = 0.1ℓ8; tℓ11 = 0.0ℓ10;
5 while (tℓ13 <ℓ15 10.0ℓ14)ℓ59 {
6   y1newℓ24 = y1ℓ17 +ℓ23 y2ℓ19 *ℓ22 hℓ21;
7   aux1ℓ28 = sin(y1ℓ26)ℓ27;
8   aux2ℓ40 = aux1ℓ30 *ℓ39 hℓ32
9     *ℓ38 gℓ34 /ℓ37 lℓ36;
10  y2newℓ46 = y2ℓ42 -ℓ45 aux2ℓ44;
11  tℓ52 = tℓ48 +ℓ51 hℓ50;
12  y1ℓ55 = y1newℓ54;
13  y2ℓ58 = y2newℓ57; };
14 require_nsb(y2,20)ℓ61;
```

$$C_1 = \begin{cases} \text{nsb}(\ell_{17}) \geq \text{nsb}(\ell_{23}) + (-1) + \xi(\ell_{23})(\ell_{17}, \ell_{22}) - (-1) // \text{ADD} \\ \text{nsb}(\ell_{22}) \geq \text{nsb}(\ell_{23}) + 0 + \xi(\ell_{23})(\ell_{17}, \ell_{22}) - (1) // \text{ADD} \\ \text{nsb}(\ell_{19}) \geq \text{nsb}(\ell_{22}) + \xi(\ell_{22})(\ell_{19}, \ell_{21}) - 1 // \text{MULT} \\ \text{nsb}(\ell_{21}) \geq \text{nsb}(\ell_{22}) + \xi(\ell_{22})(\ell_{19}, \ell_{21}) - 1 // \text{MULT} \\ \text{nsb}(\ell_{23}) \geq \text{nsb}(\ell_{24}) // \text{ASSIGN} \\ \xi(\ell_{23})(\ell_{17}, \ell_{22}) \geq 1, \xi(\ell_{22})(\ell_{19}, \ell_{21}) \geq 1 // \text{CARRY BIT} \end{cases}$$

- Constraints generated for **all the statements** of POP programs: if, while, for, array, sqrt, sin, etc.
- Linear number of constraints / variables in the size of the analyzed program

POP Label File

[Nielson, F. and Nielson, H. R., and Hankin, C. Principles of Program Analysis. Springer-Verlag, Berlin, Heidelberg'99]

```
1 g|20| = 9.81|20|; l|20| = 1.5|20|;
2 y1|29| = 0.785398|29|;
3 y2|22| = 0.0|22|;
4 h|22| = 0.1|22|; t|21| = 0.0|21|;
5 while (t<1.0) {
6   y1new|20| = y1|21| +|20| y2|22|*|22| h|22|;
7   aux1|20| = sin(y1|29|)|20|;
8   aux2|20| = aux1|20| *|20| h|20|*|20| g|20|/
          |20|l|20|;
9   y2new|20| = y2|21| -|20| aux2|18|;
10  t|20| = t|21| +|20| h|17|;
11  y1|20| = y1new|20|;
12  y2|20| = y2new|20|;
13 };
14 require_nsb(y2,20);
```

File pop_output

- ▶ 2385 bits (originally) VS **861 bits** at bit level after POP analysis

```
1 g|20| = 9.81|20|; l|20| = 1.5|20|;  
2 y1|29| = 0.785398|29|;  
3 y2|22| = 0.0|22|;  
4 h|22| = 0.1|22|; t|21| = 0.0|21|;  
5 while (t<1.0) {  
6   y1new|20| = y1|21| +|20| y2|22|*|22| h|22|;  
7   aux1|20| = sin(y1|29|)|20|;  
8   aux2|20| = aux1|20| *|20| h|20|*|20| g|20|/  
           |20|l|20|;  
9   y2new|20| = y2|21| -|20| aux2|18|;  
10  t|20| = t|21| +|20| h|17|;  
11  y1|20| = y1new|20|;  
12  y2|20| = y2new|20|;  
13 };  
14 require_nsb(y2,20);
```

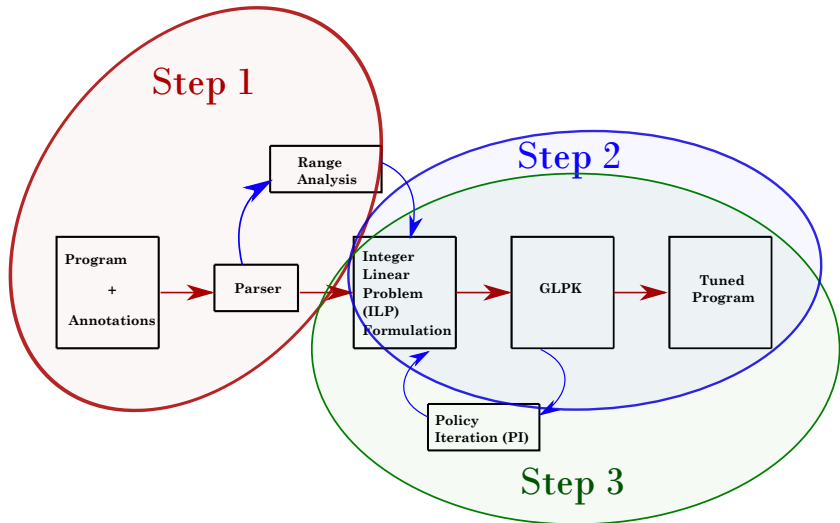
File pop_output

- ▶ 2385 bits (originally) VS 861 bits at bit level after POP analysis

How to be less pessimistic on carries propagation?

Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation




Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation



12




Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation



12




- Very **costly** over-approximation of ξ in ILP formulation

Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation



12



► Very **costly** over-approximation of ξ in ILP formulation

► $x^l = c_1^{l_1} + c_2^{l_2}$


$$\xi(l, l_1, l_2) = \begin{cases} 0 & \text{ulpe}(l_1) > \text{ufpe}(l_2) \\ 0 & \text{ulpe}(l_2) > \text{ufpe}(l_1) \\ 1 & \text{otherwise} \end{cases}$$

Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation



12



- ▶ Very **costly** over-approximation of ξ in ILP formulation

- ▶ $x^l = c_1^{l_1} + c_2^{l_2}$

$$\xi(l, l_1, l_2) = \begin{cases} 0 & \text{ulp}_e(l_1) > \text{ulp}_e(l_2) \\ 0 & \text{ulp}_e(l_2) > \text{ulp}_e(l_1) \\ 1 & \text{otherwise} \end{cases}$$

$$\xi(l)(l_1, l_2) = \min \left(\begin{array}{l} \max(\text{ufp}(l_2) - \text{ufp}(l_1) + \text{nsb}(l_1) - \text{nsb}(l_2) - \text{nsb}_e(l_2), 0), \\ \max(\text{ufp}(l_1) - \text{ufp}(l_2) + \text{nsb}(l_2) - \text{nsb}(l_1) - \text{nsb}_e(l_1), 0), 1 \end{array} \right)$$


- ▶ Need to estimate the integer quantity nsb_e

Constraints Generation for Bit-Level Tuning

Step 3: PI for Optimized Carry Bit Propagation



12



- ▶ Very **costly** over-approximation of ξ in ILP formulation

- ▶ $x^l = c_1^{l_1} + c_2^{l_2}$

$$\xi(l, l_1, l_2) = \begin{cases} 0 & \text{ulp}_e(l_1) > \text{ulp}_e(l_2) \\ 0 & \text{ulp}_e(l_2) > \text{ulp}_e(l_1) \\ 1 & \text{otherwise} \end{cases}$$

$$\xi(l)(l_1, l_2) = \min \left(\begin{array}{l} \max(\text{ufp}(l_2) - \text{ufp}(l_1) + \text{nsb}(l_1) - \text{nsb}(l_2) - \text{nsb}_e(l_2), 0), \\ \max(\text{ufp}(l_1) - \text{ufp}(l_2) + \text{nsb}(l_2) - \text{nsb}(l_1) - \text{nsb}_e(l_1), 0), 1 \end{array} \right)$$

- ▶ Need to estimate the integer quantity nsb_e

No longer ILP formulation!

Constraints Generation for Bit-Level Tuning

Policy Iteration to Refine Carry Bit Propagation



$$\xi(\ell)(l_1, l_2) = \min \left(\begin{array}{l} \max \left(\text{ufp}(\ell_2) - \text{ufp}(\ell_1) + \text{nsb}(\ell_1) - \text{nsb}(\ell_2) - \text{nsb}_e(\ell_2), 0 \right), \\ \max \left(\text{ufp}(\ell_1) - \text{ufp}(\ell_2) + \text{nsb}(\ell_2) - \text{nsb}(\ell_1) - \text{nsb}_e(\ell_1), 0 \right), 1 \end{array} \right)$$

$$\xi(\ell)(l_1, l_2) = \min \left(\begin{array}{l} \max \left(\pi_{11}, \pi_{12} \right), \\ \max \left(\pi_{21}, \pi_{22} \right), \pi_{31} \end{array} \right)$$

► Principle:

- Choose a policy π_0 by breaking the min – max equations
 - e.g. $\pi_0 = (\pi_{11}, \pi_{22}), \dots$
- Associate f^{π_0} for π_0
 - e.g. $f^{\pi_0} = (\pi_{11}, \pi_{22})$
- Solve corresponding ILP problem
- If no fixed point is reached, POP iterates until a solution is found
 - e.g. $\pi_1 = (\pi_{12}, \pi_{21}), \dots$

→ A policy cannot be selected twice in the running of the algorithm

→ The number of iterations is bounded by the number of policies

$$y1_{new}^{l_{24}} = y1^{l_{17}} +^{l_{23}} y2^{l_{19}} *^{l_{22}} h^{l_{21}}$$

$$C_2 = \left\{ \begin{array}{l} nsb_e(l_{23}) \geq nsb_e(l_{17}), \\ nsb_e(l_{23}) \geq nsb_e(l_{22}), \\ nsb(l_{23}) \geq -1 - 0 + nsb(l_{22}) - nsb(l_{17}) + nsb_e(l_{22}) + \xi(l_{23}, l_{17}, l_{22}), \\ nsb_e(l_{23}) \geq 0 - (-1) + nsb(l_{17}) - nsb(l_{22}) + nsb_e(l_{17}) + \xi(l_{23}, l_{17}, l_{22}), \\ nsb_e(l_{23}) \geq nsb_e(l_{24}), \\ nsb_e(l_{22}) \geq nsb(l_{19}) + nsb_e(l_{19}) + nsb_e(l_{21}) - 2, \\ nsb_e(l_{22}) \geq nsb(l_{21}) + nsb_e(l_{21}) + nsb_e(l_{19}) - 2, \\ \xi(l_{23})(l_{17}, l_{22}) = \min \left(\begin{array}{l} \max(0 - 6 + nsb(l_{17}) - nsb(l_{22}) - nsb_e(l_{17}), 0), \\ \max(6 - 0 + nsb(l_{22}) - nsb(l_{17}) - nsb_e(l_{22}), 0), 1 \end{array} \right) \end{array} \right\}$$

- ▶ $C = C_1 \cup C_2$: global set of constraints
- ▶ Activate optimized ξ instead of its over-approximation in pure ILP
- ▶ Optimal solution found in few seconds

```
1 g|20| = 9.81|20|; l|20| = 1.5|20|;
2 y1|29| = 0.785398|29|;
3 y2|21| = 0.0|21|;
4 h|21| = 0.1|21|; t|21| = 0.0|21|;
5 while (t<1.0) {
6   y1new|20| = y1|21| +|20| y2|21|*|22| h|21|;
7   aux1|20| = sin(y1|29|)|20|;
8   aux2|20| = aux1|19| *|20| h|18|*|19|g|17| /|18|
           l|17|;
9   y2new|20| = y2|21| -|20| aux2|18|;
10  t|20| = t|21| +|20| h|17|;
11  y1|20| = y1new|20|;
12  y2|20| = y2new|20|;
13 };
14 require_nsb(y2,20);
```

File pop_output

- ▶ 861 bits (ILP) VS 843 bits with PI method

POP(ILP)_[1, 2], POP(SMT)_[3, 4, 5, 6] and Precimonious



16

Program	Tool (LOC)	#Bits saved - Time in seconds			
		Threshold 10^{-4}	Threshold 10^{-6}	Threshold 10^{-8}	Threshold 10^{-10}
arclength	POP(ILP) (28)	2464b. - 1.8s.	2144b. - 1.5s.	1792b. - 1.7s.	1728b. - 1.8s.
	POP(SMT) (22)	1488b. - 4.7s.	1472b. - 3.04s.	864b. - 3.09s.	384b. - 2.9s.
	Precimonious (9)	576b. - 146.4s.	576b. - 156.0s.	576b. - 145.8s.	576b. - 215.0s.
simpson	POP(ILP) (14)	1344b. - 0.4s.	1152b. - 0.5s.	896b. - 0.4s.	896b. - 0.4s.
	POP(SMT) (11)	896b. - 2.9s.	896b. - 1.9s.	704b. - 1.7s.	704b. - 1.8s.
	Precimonious (10)	704b. - 208.1s.	704b. - 213.7s.	704b. - 207.5s.	704b. - 200.3s.
rotation	POP(ILP) (25)	2624b. - 0.47s.	2464b. - 0.47s.	2048b. - 0.54s.	1600b. - 0.48s.
	POP(SMT) (22)	1584b. - 1.85s.	2208b. - 1.7s.	1776b. - 1.6s.	1600b. - 1.7s.
	Precimonious (27)	2400b. - 9.53s.	2592b. - 12.2s.	2464b. - 10.7s.	2464b. - 7.4s.
accel.	POP(ILP) (18)	1776b. - 1.05s.	1728b. - 1.05s.	1248b. - 1.04s.	1152b. - 1.03s.
	POP(SMT) (15)	1488b. - 2.6s.	1440b. - 2.6s.	1056 - 2.4s.	960b. - 2.4s.
	Precimonious (0)	-	-	-	-

► Adjusting comparison criteria

- Consider **only** the variables optimized by Precimonious to estimate the quality of the optimization.
- Express all the error thresholds in base 10

[C. Rubio González, C. Nguyen, H. D. Nguyen, J. Demmel, W. Kahan, K. Sen, D. H. Bailey, C. Iancu, D. Hough, Precimonious: tuning assistant for floating-point precision, SC'13.]

- ▶ New approach for precision tuning with two variants of methods
 - ▶ Pure ILP with an over-approximation of the carry functions
 - ▶ PI technique for more precise carry bit function
- ▶ Limitation is the size of the problem accepted by the solver
- ▶ Reduce number of variables by assigning the same precision to the whole piece of code
- ▶ Technique adaptable for DNN's
- ▶ Code synthesis for the fixed-point arithmetic

Main References



- [1] Assalé Adjé, Dorra Ben Khalifa, and Matthieu Martel.
Fast and efficient bit-level precision tuning.
Submitted.
- [2] Dorra Ben Khalifa and Matthieu Martel.
A study of the floating-point tuning behaviour on the n-body problem.
Submitted.
- [3] Dorra Ben Khalifa and Matthieu Martel.
Precision tuning and internet of things.
In International Conference on Internet of Things, Embedded Systems and Communications, IINTEC 2019, pages 80–85. IEEE, 2019.
- [4] Dorra Ben Khalifa and Matthieu Martel.
Precision tuning of an accelerometer-based pedometer algorithm for iot devices.
In International Conference on Internet of Things and Intelligence System, IOTAIS 2020, pages 113–119. IEEE, 2020.
- [5] Dorra Ben Khalifa and Matthieu Martel.
An evaluation of pop performance for tuning numerical programs in floating-point arithmetic.
In International Conference on Information and Computer Technologies, ICICT 2021. IEEE, 2021.
- [6] Dorra Ben Khalifa, Matthieu Martel, and Assalé Adjé.
POP: A tuning assistant for mixed-precision floating-point computations.
In Formal Techniques for Safety-Critical Systems - 7th International Workshop, FTSCS 2019, volume 1165 of *Communications in Computer and Information Science*, pages 77–94. Springer, 2019.

Thank you! Questions?