



**HAL**  
open science

# A Data-Parallel Algorithm to Reliably Solve Systems of Nonlinear Equations

Frédéric Goualard, Alexandre Goldsztejn

► **To cite this version:**

Frédéric Goualard, Alexandre Goldsztejn. A Data-Parallel Algorithm to Reliably Solve Systems of Nonlinear Equations. Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT '08), 2008, Dunedin, New Zealand. pp.39-46, 10.1109/PDCAT.2008.26 . hal-00288207

**HAL Id: hal-00288207**

**<https://hal.science/hal-00288207>**

Submitted on 16 Jun 2008

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# A Data-Parallel Algorithm to Reliably Solve Systems of Nonlinear Equations

Frédéric Goualard<sup>¶,||</sup> and Alexandre Goldsztejn<sup>§</sup>

<sup>¶</sup>Université de Nantes, Nantes Atlantique Universités, CNRS, LINA, UMR 6241.  
2 rue de la Houssinière, BP 92208, F-44000 NANTES

<sup>§</sup>CNRS, LINA, UMR 6241. 2 rue de la Houssinière, BP 92208, F-44000 NANTES

## Abstract

Numerical methods based on interval arithmetic are efficient means to reliably solve nonlinear systems of equations. Algorithm `bc3revise` is an interval method that tightens variables' domains by enforcing a property called box consistency. It has been successfully used on difficult problems whose solving eluded traditional numerical methods. We present a new algorithm to enforce box consistency that is simpler than `bc3revise`, faster, and easily data parallelizable. A parallel implementation with Intel SSE2 SIMD instructions shows that an increase in performance of up to an order of magnitude and more is achievable.

**Keywords:** nonlinear equations, interval arithmetic, SIMD algorithm

## 1 Introduction

Interval methods [12] are numerical algorithms that use *interval arithmetic* [11] to avoid rounding error problems intrinsic to floating-point arithmetic [7]. They give enclosures of all solutions of nonlinear systems of equations with the guarantee that no solution is ever lost.

Straight interval extensions of classical numerical algorithms such as the Newton method are not well-suited to problems with many solutions or with large initial domains for the variables. To tackle these shortcomings, elaborate algorithms have been devised in the context of *Interval Constraint Programming* [1]; they are usually employed as the inner stage of a free-steering nonlinear Gauss-seidel method to exclude parts of a variable's domain that do not contain zeroes of a unidimensional equation. Domain tightening is achieved by enforcing some local consistency property. *Box consistency* [2] is one such consistency notion, which has been proved efficient in handling hard problems whose solving eluded traditional numerical methods for years [6]. It is usually enforced by Algorithm `bc3revise` [2], which combines a binary search with interval Newton steps [11] to isolate leftmost and rightmost zeroes of a unidimensional equation in the domain of a variable.

Thanks to ubiquitous Intel SSE2 SIMD instructions, it is possible to perform many interval operations at roughly the same cost as floating-point operations by computing the two bounds of the result in parallel (*basic interval vectorization*) [5]. We outline in Section 2 a novel way to do even better and to compute an interval function for two different intervals in parallel (a four times speed-up compared to “sequential” interval evaluation).

As all interval methods, Algorithm `bc3revise`—described in Section 3—can benefit from basic interval vectorization without any modification. On the other hand, it cannot take full advantage of the new arithmetic described in Section 2. Hence the introduction of Algorithm `sbc` in Section 4.1: it is a new algorithm that enforces box consistency by “shaving” domains from the left and right bounds inward. Experiments show that a sequential version of `sbc` is already faster than `bc3revise` on a set of test problems. We then describe in Section 4.2 an algorithm that exploits the potential for a high level of data parallelism in `sbc` by using SSE2 instructions to perform interval arithmetic evaluations of functions at four times the speed of a sequential code. Experiments are reported in Section 5 and show increases in performances over `bc3revise` of up to an order of magnitude and more.

---

<sup>||</sup>Corresponding author's email: `Frederic.Goualard@univ-nantes.fr`.

## 2 Interval Arithmetic and its Vectorization

Classical iterative numerical methods suffer from defects such as loss of solutions, absence of convergence, and convergence to unwanted attractors due to the use of floating-point numbers (aka *floats*). At the end of the fifties, Moore [11] popularized the use of intervals to control the errors made while computing with floats. Additionally, interval extensions of iterative numerical methods are always convergent.

In the following, we use the notations sponsored by Kearfott and others [9], where interval quantities are boldfaced.

Interval arithmetic replaces floating-point numbers by closed connected sets of the form  $I = [\underline{I}, \overline{I}] = \{a \in \mathbb{R} \mid \underline{I} \leq a \leq \overline{I}\}$  from the set  $\mathbb{I}$  of intervals, where  $\underline{I}$  and  $\overline{I}$  are floating-point numbers. In addition, each  $n$ -ary real function  $\phi$  with domain  $\mathcal{D}_\phi \subseteq \mathbb{R}^n$  is extended to an interval function  $\Phi$  with domain  $\mathcal{D}_\Phi \subseteq \mathbb{I}^n$  in such a way that the *containment principle* is verified:

$$\forall A \in \mathcal{D}_\phi, \forall I \in \mathcal{D}_\Phi: A \in I \implies \phi(A) \in \Phi(I), \quad (1)$$

as illustrated by the following example.

**Example 1** The natural interval extensions of addition and multiplication are defined by:

$$\begin{aligned} I_1 + I_2 &= [\downarrow \underline{I}_1 + \underline{I}_2 \downarrow, \uparrow \overline{I}_1 + \overline{I}_2 \uparrow] \\ I_1 \times I_2 &= [\min(\downarrow \underline{I}_1 \underline{I}_2 \downarrow, \downarrow \underline{I}_1 \overline{I}_2 \downarrow, \downarrow \overline{I}_1 \underline{I}_2 \downarrow, \downarrow \overline{I}_1 \overline{I}_2 \downarrow), \max(\uparrow \underline{I}_1 \underline{I}_2 \uparrow, \uparrow \underline{I}_1 \overline{I}_2 \uparrow, \uparrow \overline{I}_1 \underline{I}_2 \uparrow, \uparrow \overline{I}_1 \overline{I}_2 \uparrow)] \end{aligned}$$


where  $\downarrow r \downarrow$  (resp.,  $\uparrow r \uparrow$ ) is the greatest floating-point number smaller or equal (resp., the smallest floating-point number greater or equal) to  $r$ .

Then, given the real function  $f(x, y) = x \times x + y$ , we may define its natural interval extension by  $\mathbf{f}(x, y) = x \times x + y$ , and we have that, e.g.,  $\mathbf{f}([2, 3], [-1, 5]) = [3, 14]$ .

Implementations of interval arithmetic use outward rounding to enlarge the domains computed so as not to violate the containment principle (1), should some bounds be unrepresentable with floating-point numbers.

Interval addition, subtraction, multiplication, division, and integral exponentiation may be computed at roughly the same speed as their floating-point counterpart thanks to SIMD instructions, and in particular, to Intel SSE2 instructions.

Intel SSE2 instructions manipulate 128 bits registers that may be interpreted in various ways. Most notably, the registers may pack 2 double precision or 4 single precision floating-point numbers. An SSE2 operator may then compute 2 or 4 floating-point operations in parallel (see Figure 1(a)).



(a) SIMD floating-point arithmetic

(b) Two interval additions with one SSE2 instruction

Figure 1: Floating-point arithmetic and interval arithmetic in SSE2 registers

The direction of rounding for SSE2 instructions is selected independently of that of the Floating-Point Unit (FPU). An SSE2 instruction uses the same rounding for all operations performed in parallel. Nevertheless, thanks to simple floating-point properties, it is still possible to write algorithms that compute in parallel the two outward-rounded bounds of the result of interval operations. For example, we may use the property:

$$\downarrow a + b \downarrow = - \uparrow -a - b \uparrow$$

where  $a$  and  $b$  are floating-point numbers.

By storing the negation of the left bound of an interval, and by setting once and for all the rounding direction for SSE2 instructions to  $+\infty$ , the two interval additions  $[a, b] + [e, f]$  and  $[c, d] + [g, h]$  can be performed by the sole SIMD instruction depicted in Figure 1(b). All the other operators may be defined accordingly. Goualard’s paper [5] illustrates these principles for the case of basic interval vectorization (two double precision bounds computed in parallel). The algorithms to compute four bounds in parallel are new and are reported in an unpublished paper currently under review.

Armed with an interval library whose operators compute two interval operations in parallel, we may evaluate the interval extension of a function  $f$  for two different intervals for the same cost as *one floating-point* evaluation of  $f$ . In the following, we note  $\llbracket f(I_1), f(I_2) \rrbracket$  such a parallel evaluation of  $f$  for two different interval arguments.

### 3 Box Consistency and the bc3revise Algorithm

*Interval Constraint Programming* [1] is a successful approach to reliably isolate all solutions of systems of equations. It makes cooperate *contracting operators* to prune the domains of the variables (intervals with floating-point bounds from the set  $\mathbb{I}$ ) with smart *propagation algorithms* [10]—akin to free-steering nonlinear Gauss-Seidel—to ensure consistency among all the constraints.

The amount of pruning obtained from one equation is controlled by the level of consistency enforced. Box consistency [2] is defined as follows:

**Definition 1 (Box consistency)** *An equation  $f(x_1, \dots, x_n) = 0$  is box consistent with respect to a variable  $x_i$  and a box  $B = I_1 \times \dots \times I_n$  if and only if:*

$$\begin{cases} 0 \in f(I_1, \dots, I_{i-1}, [\underline{I}_i, \underline{I}_i^+], I_{i+1}, \dots, I_n) \\ \text{and} \\ 0 \in f(I_1, \dots, I_{i-1}, [\overline{I}_i^-, \overline{I}_i], I_{i+1}, \dots, I_n), \end{cases} \quad (2)$$

where  $I = [\underline{I}, \overline{I}]$  is an interval with floating-point bounds,  $a^+$  (resp.,  $a^-$ ) is the smallest floating-point number greater than (resp., the largest floating-point number smaller than) the floating-point number  $a$ , and  $f$  is the natural interval extension of  $f$ .

Given a real function  $f: \mathbb{R}^n \rightarrow \mathbb{R}$ , and a box of domains  $B = I_1 \times \dots \times I_n \in \mathbb{I}^n$ , we define  $g_i^B: \mathbb{I} \rightarrow \mathbb{I}$  as the  $i$ th unary interval projection with respect to  $B$  of its interval extension  $f$ :

$$g_i^B(x) = f(I_1, \dots, I_{i-1}, x, I_{i+1}, \dots, I_n), \quad i \in \{1, \dots, n\}.$$

In the following, we will mostly manipulate  $g_i^B$  instead of  $f$ . The original real function  $f$ , the box  $B$  of domains considered and/or the variable  $x_i$  on which the projection is performed will often be left implicit and omitted from the notation of  $g$ .

In order not to lose any potential solution, an algorithm that enforces box consistency of an equation with respect to a variable  $x_i$  and a box of domains must return the largest domain  $I'_i \subseteq I_i$  that verifies Eq. (2).

---

**Algorithm 1** Computing a box consistent interval with respect to  $g$  the usual way

---

[bc3revise] in:  $g: \mathbb{I} \rightarrow \mathbb{I}$ ; in:  $I \in \mathbb{I}$

**# Returns the largest interval included in  $I$  that is box consistent with respect to  $g$**

begin

1  $I_l \leftarrow \text{left\_narrow}(g, I)$

2 if  $I_l \neq \emptyset$ :

3  $I_r \leftarrow \text{right\_narrow}(g, [\underline{I}_l, \overline{I}])$

4 return  $\square(I_l \cup I_r)$  **# returns the smallest interval w.r.t. set inclusion that contains  $I_l \cup I_r$**

5 else:

6 return  $\emptyset$

end

---

Algorithm bc3revise [2], presented by Algorithms 1 and 2, considers the unary projection of an  $n$ -ary equation on a variable  $x_i$  (where all variables but  $x_i$  have been replaced by their current domain) and a domain  $I_i$ . It enforces

box consistency by searching the leftmost and rightmost *canonical domains*\* for which  $g$  evaluates to an interval containing 0. The search is performed by a dichotomic search aided with Newton steps to accelerate the process. Algorithm 2 describes the search of a quasi-zero to update the left side of  $I_i$ . The procedure `right_narrow` to update the right bound works along the same lines and is, therefore, omitted.

---

**Algorithm 2** Computing a box consistent left bound with Newton steps and a binary search

---

```

[left_narrow] in:  $g: \mathbb{I} \rightarrow \mathbb{I}$ ; in:  $I \in \mathbb{I}$ 
# Returns an interval included in  $I$ 
# with the smallest left bound  $l$  such that  $0 \in g([l, l^+])$ 
begin
  1 if  $0 \notin g(I)$ : # No solution in  $I$ 
  2   return  $\emptyset$ 
  3 else:
  4   if  $\underline{I}^+ \geq \bar{I}$ : # canonical( $I$ ):
  5     return  $I$ 
  6   else:
  7      $I \leftarrow \text{Newton}(g, g', I)$  # Interval Newton steps
  8     if  $0 \in g([\underline{I}, \underline{I}^+])$ : # Box consistent left bound?
  9       return  $I$ 
 10    else:
 11       $(I_1, I_2) \leftarrow \text{split}(I)$  #  $I_1 \leftarrow [\underline{I}, m(I)]$ ,  $I_2 \leftarrow [m(I), \bar{I}]$ 
 12       $I \leftarrow \text{left\_narrow}(g, I_1)$ 
 13      if  $I = \emptyset$ :
 14        return  $\text{left\_narrow}(g, I_2)$ 
 15      else:
 16        return  $I$ 
end

```

---

Algorithm `bc3revise` first tries to move the left bound of  $I_i$  to the right, and then proceeds to move its right bound to the left. The `Newton` procedure computes a fixpoint of the Interval Newton algorithm [11], where a Newton step at iteration  $j + 1$  is:

$$I^{(j+1)} \leftarrow I^{(j)} \cap \left( m(I^{(j)}) - \frac{g(m(I^{(j)}))}{g'(I^{(j)})} \right)$$

with  $m(I)$  the midpoint of the interval  $I^\dagger$ . As in Ratz's work [13], the Newton step uses an extended version of the interval division to return a union of two semi-open-ended intervals whenever  $g'(I^{(j)})$  contains 0. The subtraction and the intersection operators are modified accordingly. The intersection operator is applied to an interval ( $I^{(j)}$ ) and a union of two intervals (result of the subtraction), and returns an interval. Figure 2 presents graphically the steps performed to enforce box consistency. The encircled numbers label the steps.

Algorithms `bc3revise`, `left_narrow` and `right_narrow` do not offer opportunities to exploit full data parallelism as they do not require close evaluations of the same function over different domains. The same holds for the `Newton` procedure: in the general case,  $g$  and  $g'$  are different functions, and therefore cannot be evaluated in parallel with SIMD instructions.

## 4 Box Consistency by Shaving

To obtain a higher level of data parallelism, we propose a new algorithm to enforce box consistency on the projection  $g_i^B$  of an equation  $f = 0$  on a variable  $x_i$ : starting from the original domain  $I_i$  for  $x_i$ , consider separately its left half and its right half; for the left part, linearize  $g$  at  $\underline{I}_i$  as for a Newton step, solve the resulting linear equation and intersect the resulting domain with the left half of  $I_i$ ; do the same for the right half by linearizing  $g$  at  $\bar{I}_i$ . A new smaller domain that preserves solutions is then obtained; Iterate until reaching a fixpoint.

---

\* A non-empty interval  $[a, b]$  is canonical if  $a^+ \geq b$ .

† Note that we are free to choose any point in  $I$ , not the midpoint only. We take advantage of this in the algorithms presented in the next section.




Figure 2: Enforcing box consistency with bc3revise

#### 4.1 A Sequential Algorithm: sbc

Figure 3 illustrates graphically the process just described, and Algorithm 3 presents the actual algorithm.

**Proposition 1 (Termination, Correctness, and Completeness of sbc)** *Given an  $n$ -ary equation  $c: f(x_1, \dots, x_n) = 0$ , a box  $\mathbf{I}_1 \times \dots \times \mathbf{I}_n$  of domains, and a projection  $c: g_i(x) = 0$  of  $c$ , we have:*

**Termination.** *The call to  $\text{sbc}(g_i, \mathbf{I}_i)$  always terminates;*

**Correctness.** *The equation  $c$  is box consistent with respect to  $x_i$  and  $\text{sbc}(g_i, \mathbf{I}_i)$ ;*

**Completeness.** *No solution is lost during the tightening process:*

$$\forall (r_1, \dots, r_n) \in \mathbf{I}_1 \times \dots \times \mathbf{I}_n: f(r_1, \dots, r_n) = 0 \implies r_i \in \text{sbc}(g_i, \mathbf{I}_i).$$

**Proof.** *In the following, the interval  $\mathbf{I}$  corresponds to the domain  $\mathbf{I}_i$  of  $x_i$ .*

*(Termination).* Algorithm sbc terminates in any case since we always tighten  $\mathbf{I}$  in the loop 2–23 (either by splitting it on Line 3, or by tightening  $\mathbf{I}_l$  and  $\mathbf{I}_r$  with Newton steps on Lines 10 and 19). Since we consider intervals with floating-point bounds, of which there are finitely many, we have to reach canonicity of  $\mathbf{I}$  eventually. At that point, either  $g_i(\mathbf{I})$  contains 0, and we have reached box consistency, or it does not, and we can safely narrow  $\mathbf{I}$  to  $\emptyset$ , which both make us leave the loop.

*(Correctness).* We leave the loop 2–23 if  $\mathbf{I}$  is empty or if the canonical intervals at left and right bounds contain solutions of  $g_i$  (Lines 5 and 14). In the latter case, we have the two conditions of Eq (2) for box consistency; the former case occurs if both  $\mathbf{I}_l$  and  $\mathbf{I}_r$  do not contain a solution of  $c$  (Lines 7 and 16) or if the Newton steps on Lines 10 and 19 narrow them down to  $\emptyset$ . By correctness of the interval Newton method, this case only occurs if, once again,  $\mathbf{I}_l$  and  $\mathbf{I}_r$  do not contain a solution of  $c$ .

*(Completeness).* By completeness of the Newton operator, as tightening only occurs either through Newton steps, or by discarding intervals that have been proved on Lines 5, 7, 14, or 16 not to contain solutions.  $\square$

For each iteration of the loop 2–23 in Algorithm sbc, we have to compute  $g$  for intervals  $[\underline{\mathbf{I}}_l, \underline{\mathbf{I}}_l^+]$ ,  $[\overline{\mathbf{I}}_r^-, \overline{\mathbf{I}}_r]$ ,  $\mathbf{I}_l$ ,  $\mathbf{I}_r$ ,  $[\underline{\mathbf{I}}_l, \underline{\mathbf{I}}_l]$ , and  $[\overline{\mathbf{I}}_r, \overline{\mathbf{I}}_r]$ . We also have to compute  $g'$  for intervals  $\mathbf{I}_l$  and  $\mathbf{I}_r$ . All these evaluations are candidates for parallelization with SIMD instructions, as presented in the next section.

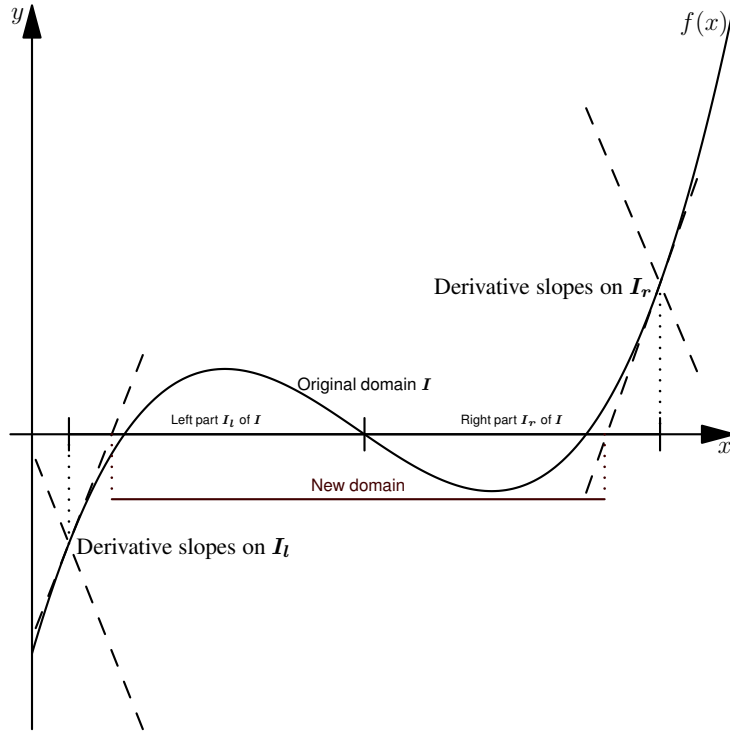


Figure 3: Domain reduction obtained with one iteration of the loop in sbc

## 4.2 An SIMD Algorithm for Box Consistency: vsbc

Algorithm 4 presents a modification of Algorithm sbc to make good use of its higher level of data parallelism thanks to the SIMD interval arithmetic that has been presented in Section 2. The evaluations of  $g$  and  $g'$  are reordered to appear in pairs that can be evaluated in parallel. In addition, we reuse the evaluation of  $g([\underline{I}_l, \underline{I}_l^+])$  and  $g([\overline{I}_r^-, \overline{I}_r])$  of Line 4 for the Newton steps of Line 22 instead of  $g([\underline{I}_l, \underline{I}_l])$  and  $g([\overline{I}_r, \overline{I}_r])$  as was done in Algorithm sbc. This choice avoids two evaluations of  $g$  at the cost of potentially slightly decreasing the tightening ability of the Newton step. The domain computed is unaffected by this optimization. In particular, box consistency is still obtained eventually.

At each iteration of the loop between Lines 2 and 24, we perform 4 *interval* evaluations of  $g$  and 2 *interval* evaluations of  $g'$  for the same cost as 2 *floating-point* evaluations of  $f$  and 1 *floating-point* evaluation of  $f'$ .

## 5 Experiments

To evaluate the impact of our new algorithms, we have selected 20 instances of 12 classical test problems. Some are polynomial and others are not. The characteristics of these test problems are summarized in Table 1. All problems are structurally well constrained, with as many equations as variables. Column “Size” reports the number of equations/variables. Column “Equations” indicates whether all equations are polynomial (quadratic, if no polynomial has a degree greater than 2). A problem is labelled “non-polynomial” if at least one constraint contains a trigonometric, hyperbolic or otherwise transcendental operator. All test problems are presented on the COPRIN web page [8].

All experiments were conducted on an Intel Core2 Duo T5600 1.83GHz. The Whetstone test [3] for this machine reports 1111 MIPS with a loop count equal to 100,000. All algorithms have been implemented in an in-house C++ solver, with *gaol* [4] as the underlying interval arithmetic library. The SIMD interval arithmetic presented above has been implemented from scratch using Intel intrinsic instructions. In its current state, the library only contains vectorized versions of the addition, subtraction, multiplication, division, and integral power. All other SIMD functions are emulated with sequential interval arithmetic. As a consequence, only polynomial equation systems are entirely solved in an SIMD environment at present.

Table 2 reports the time spent in seconds to isolate all solutions of the test problems in domains with a width



---

**Algorithm 3** Enforcing box consistency by shaving

---

[**sbc**] in:  $g: \mathbb{I} \rightarrow \mathbb{I}$ ; in:  $I \in \mathbb{I}$

# Returns the largest interval included in  $I$  that is box consistent w.r.t.  $g$

begin

```
1 (left_consistent, right_consistent)  $\leftarrow$  (false, false)
2 do:
3   ( $I_l$ ,  $I_r$ )  $\leftarrow$  split( $I$ )
4   if  $\neg$ left_consistent: # Updating the left bound
5     if  $0 \notin g([I_l, I_l^+])$ : #  $I$  not box consistent to the left?
6        $I_l \leftarrow [I_l^+, \bar{I}_l]$  # Considering the remainder of  $I_l$ 
7       if  $0 \notin g(I_l)$ : # No solution in the remainder of  $I_l$ ?
8          $I_l \leftarrow \emptyset$ 
9       else:
10         $I_l \leftarrow I_l \cap (I_l - g([I_l, I_l])/g'(I_l))$  # One Newton step
11      else:
12        left_consistent  $\leftarrow$  true
13    if  $\neg$ right_consistent: # Updating the right bound
14      if  $0 \notin g([\bar{I}_r^-, I_r])$ : #  $I$  not box consistent to the right?
15         $I_r \leftarrow [I_r, \bar{I}_r^-]$  # Considering the remainder of  $I_r$ 
16        if  $0 \notin g(I_r)$ : # No solution in the remainder of  $I_r$ ?
17           $I_r \leftarrow \emptyset$ 
18        else:
19           $I_r \leftarrow I_r \cap (\bar{I}_r - g([\bar{I}_r, \bar{I}_r])/g'(I_r))$  # One Newton step
20      else:
21        right_consistent  $\leftarrow$  true
22     $I \leftarrow \square(I_l \cup I_r)$  # Returns the "hull"  $[\min(I_l, I_r), \max(\bar{I}_l, \bar{I}_r)]$  of the union
23    while  $(I \neq \emptyset) \wedge (\neg$ left_consistent  $\vee \neg$ right_consistent)
24    return  $I$ 
```

end

---

smaller than  $10^{-6}$ , starting from the standard domains given on the COPRIN web page. An entry "TO" indicates a time-out (more than 30 minutes, here). Column **bc3revise** presents the results obtained with Algorithm **bc3revise** implemented with double precision interval arithmetic on the FPU; Column **bc3vd** corresponds to Algorithm **bc3revise** where interval arithmetic is performed in double precision with SSE2 instructions (basic vectorization); Column **bc3vf** corresponds to Algorithm **bc3revise** where interval arithmetic is performed in single precision with SSE2 instructions (we still perform only one interval operation per SSE2 instruction, using only the lower half of SSE2 registers, though); Column **sbc** corresponds to Algorithm **sbc** implemented with double precision interval arithmetic on the FPU; Column **sbcvd** corresponds to Algorithm **sbc** where interval arithmetic is performed in double precision with SSE2 instructions; Column **vsbc** corresponds to Algorithm **vsbc** where interval arithmetic is performed in single precision with SSE2 instructions (two interval operations are performed in parallel); lastly, Column **vsbc&sbcvd** corresponds to the cooperation of **vsbc** and **sbcvd**: **vsbc** is used until the domains are all smaller than a size fixed empirically to 0.25; **sbcvd** is used afterwards.

## 6 Discussion

As can be seen from Column **sbc** of Table 2, enforcing box consistency by shaving is faster than with Algorithm **bc3revise** on all problems of our test set, the ratio **bc3revise/sbc** ranging from 1.9 to 17.9 and beyond. We also believe that **sbc** is simpler to understand and easier to implement correctly than **bc3revise**.



---

**Algorithm 4** A data parallel algorithm for box consistency enforcement

---

[vsbc] in:  $g: \mathbb{I} \rightarrow \mathbb{I}$ ; in:  $I \in \mathbb{I}$

# Returns the largest interval included in  $I$  that is box consistent w.r.t.  $g$

begin

```
1 (left_consistent, right_consistent) ← (false, false)
2 do:
3   ( $I_l, I_r$ ) ← split( $I$ )
4   ( $J_l, J_r$ ) ←  $\llbracket g(\underline{I}_l, \underline{I}_l^+), g(\overline{I}_r^-, \overline{I}_r) \rrbracket$ 
5   if  $0 \notin J_l$ : #  $I$  not box consistent to the left?
6      $I_l \leftarrow [\underline{I}_l^+, \overline{I}_l]$  # Considering the remainder of  $I_l$ 
7   else:
8     left_consistent ← true
9   if  $0 \notin J_r$ : #  $I$  not box consistent to the right?
10     $I_r \leftarrow [\underline{I}_r, \overline{I}_r^-]$  # Considering the remainder of  $I_r$ 
11  else:
12    right_consistent ← true
13  if  $\neg$ left_consistent  $\vee$   $\neg$ right_consistent:
14    ( $K_l, K_r$ ) ←  $\llbracket g(I_l), g(I_r) \rrbracket$ 
15    if  $0 \notin K_l$ : # First checking an obvious absence of solution in  $I_l$ 
16       $I_l \leftarrow \emptyset$ 
17    if  $0 \notin K_r$ : # First checking an obvious absence of solution in  $I_r$ 
18       $I_r \leftarrow \emptyset$ 
19    # Performing 2 Newton steps in parallel to update both bounds
20    # For better performances, we reuse  $J_l$  and  $J_r$ 
21    # instead of  $g([\underline{I}_l, \underline{I}_l])$  and  $g([\overline{I}_r, \overline{I}_r])$ 
22    ( $I_l, I_r$ ) ←  $\llbracket I_l \cap ([\underline{I}_l, \underline{I}_l^+] - J_l/g'(I_l)), I_r \cap ([\overline{I}_r^-, \overline{I}_r] - J_r/g'(I_r)) \rrbracket$ 
23     $I \leftarrow \square(I_l \cup I_r)$  # Returns the "hull"  $[\min(\underline{I}_l, \underline{I}_r), \max(\overline{I}_l, \overline{I}_r)]$  of the union
24  while ( $I \neq \emptyset$ )  $\wedge$  ( $\neg$ left_consistent  $\vee$   $\neg$ right_consistent)
25  return  $I$ 
```

end

---

Table 1: Test problems characteristics

Name	Code	Size	Equations
Bronstein	<i>bro</i>	3	quadratic
Broyden-banded	<i>bb</i>	100, 500, and 1 000	quadratic
Broyden tridiagonal	<i>bt</i>	10 and 20	quadratic
Combustion	<i>comb</i>	10	polynomial
Discrete Boundary Value Function	<i>dbvf</i>	10 and 30	polynomial
Extended Freudenstein	<i>ef</i>	30 and 50	polynomial
Mixed Algebraic Trigonometric	<i>mat</i>	3	non-polynomial
Moré-Cosnard	<i>mc</i>	50 and 100	polynomial
Robot	<i>rob</i>	8	quadratic
Trigexp 3	<i>te3</i>	5 000	non-polynomial
Troesch	<i>tro</i>	50, 100, and 200	non-polynomial
Yamamura	<i>yam</i>	6 and 8	polynomial

Table 2: Experiments

Problem	bc3revise	bc3vd	bc3vf	sbc	sbcvd	vsbc	vsbcvd&sbcvd
<i>bro</i>	2.6	0.9	0.4	0.3	0.3	<b>0.06</b>	0.2
<i>bb</i> 100	10.4	3.3	3.2	2.7	1.2	<b>0.7</b>	0.8
<i>bb</i> 500	123.7	39.1	27.0	24.8	10.8	<b>5.3</b>	5.5
<i>bb</i> 1000	280.2	88.7	56.8	55.1	24.0	<b>11.1</b>	11.4
<i>bt</i> 10	16.9	5.8	3.4	2.1	1.0	<b>0.4</b>	0.5
<i>bt</i> 20	1127.4	382.2	260.3	141.1	66.7	<b>28.6</b>	30.4
<i>comb</i>	1.4	0.6	0.5	0.2	0.08	<b>0.03</b>	0.08
<i>dbvf</i> 10	1.7	0.5	0.6	0.3	0.1	<b>0.08</b>	0.1
<i>dbvf</i> 30	42.7	13.4	20.4	7.7	<b>3.2</b>	4.7	3.8 $\nabla$
<i>ef</i> 30	2.1	0.8	<b>TO</b>	1.1	0.5	<b>TO</b>	<b>0.3</b>
<i>ef</i> 50	5.1	1.8	<b>TO</b>	2.2	0.9	<b>TO</b>	<b>0.6</b>
<i>mat</i>	26.9	13.6	13.2	1.5	1.1	<b>0.4</b>	1.0
<i>mc</i> 50	23.7	6.5	5.7	11.0	4.4	<b>3.0</b>	4.1
<i>mc</i> 100	175.5	49.0	46.2	86.9	36.8	<b>28.0</b>	35.6
<i>rob</i>	4.5	1.5	4.3	0.8	0.4	1.0	<b>0.3</b>
<i>te3</i> 5000	7.5	5.1	17.1	2.9	<b>2.7</b>	3.3	3.8 $\nabla$
<i>tro</i> 50	24.8	15.3	29.1	4.4	3.6	<b>2.6</b>	3.4
<i>tro</i> 100	180.1	112.7	384.2	30.9	25.4	33.7	<b>24.2</b>
<i>tro</i> 200	1341.4	844.4	<b>TO</b>	231.1	188.1	<b>TO</b>	<b>181</b>
<i>yam</i> 6	14.6	4.3	4.2	7.3	2.5	<b>1.2</b>	1.7
<i>yam</i> 8	279.1	84.6	104.0	91.0	35.1	<b>26.1</b>	27.7

Times in seconds on an Intel Core2 Duo T5600 1.83GHz (whetstone 100 000=1111 MIPS)

Best times in bold blue. Time out **TO** set to 1 800 s.

Numbers followed by the symbol " $\nabla$ " correspond to cases for which vsbcvd&sbc performs worse than sbc alone.

Basic vectorization of interval arithmetic improves speed by up to three times (see bc3revise vs. bc3vd and sbc vs. sbcvd) at no cost since algorithms do not have to be modified in order to benefit from it.

If we take advantage of the data parallelism inherent to Algorithm sbc to vectorize interval evaluations, leading to Algorithm vsbc, we obtain even better results on all problems but *dbvf* 30, *ef* 30, *ef* 50, *rob*, *te3* 5000, *tro* 100, and *tro* 200. All other things being equal, if we emulate SIMD instructions with double precision floating-point operations, we obtain back the times of sbc, which means that the very size of single precision floating-point numbers is the culprit here: as we vectorize 2 interval instructions with SSE2 registers, we must switch from double precision floats used in the rest of the program to single precision floats (see Figure 1(b)). The cast leads to less precision in the computation, which in turn has an impact on the ability to reject domains having no solutions. The same problem occurs for bc3vf.




Figure 4: Slow convergence with single precision floating-point numbers

As a consequence, the exploration algorithm has more branching to perform to isolate solutions. This incurs an increase in the running time that may be drastic for ill-conditioned problems such as *Troesch* or *Extended Freudenstein*: as we may see in Figure 4 for the case of 2 equations and 2 variables, the curves for these two problems are almost tangent to each other and to the  $xy$ -plane on a large surface. Each equation considered separately leads to the computation of many quasi-zeroes that cannot be removed easily by the other equation of the problem.

There is currently no easy cure to this problem as microprocessor makers do not seem to be ready to offer SIMD instructions on 4 double precision floats any time soon. It is still possible to quickly isolate regions of interest in “large” domains using `vsbc`, and then switch to `sbcvd` to polish the results and obtain tighter domains if necessary. Column `vsbc&sbcvd` in Table 2 shows that this procedure indeed removes the time-out problems of `vsbc` on ill-conditioned problems, while still preserving better performances compared to `sbcvd` alone. The best cooperation scheme that maximizes performances as much as possible still remains to be found, though.

## References

- [1] Frédéric Benhamou. Interval constraints, interval propagation. In Panos M. Pardalos and Christodoulos A. A. Floudas, editors, *Encyclopedia of Optimization*, volume 3, pages 45–48. Kluwer Academic Publishers, Dordrecht, The Netherlands, 2001.
- [2] Frédéric Benhamou, David McAllester, and Pascal Van Hentenryck. CLP(Intervals) revisited. In *Procs. Intl. Symp. on Logic Prog.*, pages 124–138. The MIT Press, 1994.
- [3] H. J. Curnow and Brian A. Wichmann. A synthetic benchmark. *Comput. J.*, 19(1):43–49, 1976.
- [4] Frédéric Goualard. *GAOL 3.1.1: Not Just Another Interval Arithmetic Library*. Laboratoire d’Informatique de Nantes-Atlantique, 4.0 edition, October 2006. Available at <http://sourceforge.net/projects/gaol>.
- [5] Frédéric Goualard. Fast and correct SIMD algorithms for interval arithmetic. In *Proceedings of PARA ’08*, May 2008.
- [6] Laurent Granvilliers and Frédéric Benhamou. Progress in the solving of a circuit design problem. *Journal of Global Optimization*, 20:155–168, 2001.
- [7] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
- [8] INRIA project COPRIN: Contraintes, OPTimisation, Résolution par INtervalles. The COPRIN examples page. Web page at <http://www-sop.inria.fr/coprin/logiciels/ALIAS/Benches/benches.html>.
- [9] Ralph Baker Kearfott, Mitsuhiro T. Nakao, Arnold Neumaier, Siegfried M. Rump, Sergey P. Shary, and Pascal van Hentenryck. Standardized notation in interval analysis. In *Proc. XIII Baikal International School-seminar “Optimization methods and their applications”*, volume 4 “Interval analysis”, pages 106–113. Irkutsk: Institute of Energy Systems SB RAS, July 2005.
- [10] Alan K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 1(8):99–118, 1977.
- [11] Ramon Edgar Moore. *Interval Analysis*. Prentice-Hall, Englewood Cliffs, N. J., 1966.
- [12] Arnold Neumaier. *Interval methods for systems of equations*, volume 37 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1990.
- [13] Dietmar Ratz. Inclusion isotone extended interval arithmetic. Technical Report 5/1996, Institut für Angewandte Mathematik, Universität Karlsruhe, 1996.