

# Parallel parametric linear programming solving, and application to polyhedral computations

Camille Coti<sup>1</sup>, David Monniaux<sup>2</sup>, and Hang Yu<sup>2</sup>

<sup>1</sup> LIPN, CNRS UMR 7030, Université Paris 13, Sorbonne Paris Cité  
99, avenue Jean-Baptiste Clément, F-93430 Villetaneuse, FRANCE  
`Camille.Coti@lipn.univ-paris13.fr`

<sup>2</sup> Univ. Grenoble Alpes, CNRS, Grenoble INP\*, F-38000 Grenoble, France  
`First-name.Last-name@univ-grenoble-alpes.fr`

**Abstract.** Parametric linear programming is central in polyhedral computations and in certain control applications. We propose a task-based scheme for parallelizing it, with quasi-linear speedup over large problems.

## 1 Introduction

A *convex polyhedron*, or *polyhedron* for short here, in dimension  $n$  is the solution set over  $\mathbb{Q}^n$  (or, equivalently,  $\mathbb{R}^n$ ) of a system of inequalities (with integer or rational coefficients). Polyhedra in higher dimension are typically used to enclose the reachable states of systems whose state can be expressed, at least partially, as a vector of reals or rationals; e.g. hybrid systems or software [3].

The conventional approaches for polyhedral computations are the *dual description* (using both vertices and faces) and *Fourier-Motzkin elimination*. They both suffer from high complexity on relevant cases. We instead express image, projection, convex hull etc. as solutions to *parametric linear programmings*, where parameters occur linearly within the objective. A solution to such a program is a quasi-partition of the space of parameters into polyhedra, with one optimum associated to each polyhedron. The issue is how to compute this solution efficiently. In this article, we describe how we parallelized our algorithm.

## 2 Sequential algorithms

Here we are leaving out how polyhedral computations such as projection and convex hull can be reduced to parametric linear programming — this is covered in the literature [4,7] — and focus on solving the parametric linear programs.

### 2.1 Non-parametric linear programming (LP)

A linear program with  $n$  unknowns is defined by a system of equations  $AX = B$ , where  $A$  is an  $m \times n$  matrix; a solution is a vector  $X$  such that  $X \geq 0$  on all

---

\* Institute of Engineering Univ. Grenoble Alpes

coordinates and  $AX = B$ . The program is said to be *feasible* if it has at least one solution, *infeasible* otherwise. In a non-parametric linear program one considers an objective  $C$ : one wants the solution that maximizes  $C^T X$ . The program is deemed *unbounded* if it is feasible yet it has no such optimal solution.

*Example 1.* Consider the polygon  $P$  defined by  $x_1 \geq 0$ ,  $x_2 \geq 0$ ,  $3x_1 - x_2 \leq 6$ ,  $-x_1 + 3x_2 \leq 6$ . Define  $x_3 = 6 - 3x_1 + x_2$  and  $x_4 = 6 + x_1 - 3x_2$ . Let  $X = (x_1, x_2, x_3, x_4)$ , and then  $P$  is the projection onto the first two coordinates of the solution set of  $AX = B \wedge X \geq 0$  where  $A = \begin{bmatrix} 1 & -3 & 0 & -1 \\ -3 & 1 & -1 & 0 \end{bmatrix}$  and  $B = \begin{bmatrix} 6 \\ 6 \end{bmatrix}$ .

An LP solver takes as input  $(A, B, C)$  and outputs “infeasible”, “unbounded” or an optimal solution. Most solvers work with floating-point numbers and their final answer may be incorrect: they may answer “infeasible” whereas the problem is feasible, or give “optimal solutions” that are not solutions, or not optimal.

In addition to a solution  $X^*$ , solvers also provide the associated *basis*:  $X^*$  is defined by setting  $n - m$  of its coordinates to 0 (known as *nonbasic variables*) and solving for the other coordinates (known as *basic variables*) using  $AX^* = B$ , and the solver provides the partition into basic and nonbasic variables it used. If a floating-point solver is used, it is possible to reconstruct an exact rational point  $X^*$  using that information and a library for solving linear systems in rational arithmetic. One then checks whether it is truly a solution by checking  $X^* \geq 0$ .

The optimal basis also contains a proof of optimality of the solution. We compute the objective function  $C^T X$  as  $\sum_{i \in N} \alpha_i X_i + c$  where  $N$  is the set of indices of the nonbasic variables and  $c$  is a constant, and conclude that the solution obtained by setting these nonbasic variables to 0 is maximal because all the  $\alpha_i$  are nonpositive. If  $X^*$  is not a solution of the problem ( $X^* \geq 0$  fails) or is not optimal, then we fall back to an exact implementation of the simplex algorithm.

*Example 1 (continued).* Assume the objective is  $C = [1 \ 1 \ 0 \ 0]$ , that is,  $C^T X = x_1 + x_2$ . From  $AX = B$  we deduce  $x_1 = 3 - \frac{3}{8}x_3 - \frac{1}{8}x_4$  and  $x_2 = 3 - \frac{1}{8}x_3 - \frac{3}{8}x_4$ . Thus  $x_1 + x_2 = 6 - \frac{1}{2}x_3 - \frac{1}{2}x_4$ .

Assume  $x_3$  and  $x_4$  are nonbasic variables and thus set to 0, then  $X^* = (x_1, x_2, x_3, x_4) = (3, 3, 0, 0)$ . It is impossible to improve upon this solution: as  $X \geq 0$ , changing the values of  $x_3$  and  $x_4$  can only decrease the objective  $o = 6 - \frac{1}{2}x_3 - \frac{1}{2}x_4$ . This expression of  $o$  from the nonbasic variables can be obtained by linear algebra once the partition into basic and nonbasic variables is known.

While the optimal value  $C^T X^*$ , if it exists, is unique for a given  $(A, B, C)$ , there may exist several  $X^*$  for it, a situation known as *dual degeneracy*. The same  $X^*$  may be described by different bases, a situation known as *primal degeneracy*, happening when more than  $n - m$  coordinates of  $X^*$  are zero, and thus some basic variables could be used as nonbasic and the converse.

## 2.2 Parametric linear programming (PLP)

For a *parametric* linear program, we replace the constant vector  $C$  by  $C_0 + \sum_{i=1}^k \mu_i C_i$  where the  $\mu_i$  are parameters.<sup>3</sup> When the  $\mu_i$  change, the optimum  $X^*$  changes. Assume temporarily that there is no degeneracy. Then, for given values of the  $\mu_i$ , the problem is either unbounded, or there is one single optimal solution  $X^*$ . It can be shown that the region of the  $(\mu_1, \dots, \mu_k)$  associated to a given optimum  $X^*$  is a convex polyhedron (if  $C_0 = 0$ , a convex polyhedral cone), and that these regions form a quasi partition of the space of parameters (two regions may overlap at their boundary, but not in their interior) [4,5,7]. The output of the parametric linear programming solver is this quasi-partition, and the associated optima—in our applications, the problem is always bounded in the optimization directions, so we do not deal with the unbounded case.

Let us see in more details about how to compute these regions. We wish to attach to each basis (at least, each basis that is optimal for at least one vector of parameters) the region of parameters for which it is optimal.

*Example 1 (continued).* Instead of  $C = [1 \ 1 \ 0 \ 0]$  we consider  $C = [\mu_1 \ \mu_2 \ 0 \ 0]$ . Let us now express  $o = C^T X$  as a function of the nonbasic variables  $x_3$  and  $x_4$ :

$$o = (3\mu_1 + 3\mu_2) + \left(-\frac{3}{8}\mu_1 - \frac{1}{8}\mu_2\right)x_3 + \left(-\frac{1}{8}\mu_1 - \frac{3}{8}\mu_2\right)x_4 \quad (1)$$

The coefficients of  $x_3$  and  $x_4$  are nonpositive if and only if  $3\mu_1 + \mu_2 \geq 0$  and  $\mu_1 + 3\mu_2 \geq 0$ , which define the cone of optimality associated to that basis and to the optimum  $X^* = (3, 3, 0, 0)$ .

The description of the optimality polyhedron by the constraints obtained from the sign conditions in the objective function may be redundant: containing constraints that can be removed without changing the polyhedron. Our procedure [6] for removing redundant constraints from the description of a region  $R_1$  also provides a set of vectors outside of  $R_1$ , a feature that will be useful.

Assume now we have solved the optimization problem for a vector of parameters  $D_1$ , and obtained a region  $R_1$  in the parameters (of course,  $D_1 \in R_1$ ). We store the set of vectors outside of  $R_1$  provided by the redundancy elimination procedure into a “working set”  $W$  to be processed, choose  $D_2$  in it. We compute the region  $R_2$  associated to  $D_2$ . Assume that  $R_2$  and  $R_1$  are adjacent, meaning that they have a common boundary. We get vectors outside of  $R_2$  and add them to  $W$ . We pick  $D_3$  in  $W$ , check that it is not covered by  $R_1$  or  $R_2$ , and, if it is not, compute  $R_3$ , etc. The algorithm terminates when  $W$  becomes empty, meaning the  $R_1, \dots$  produced form the sought quasi-partition.

This simplistic algorithm can fail to work because it assumes that it is discovering the adjacency relation of the graph. The problem is that, if we move from a region  $R_i$  to a vector  $D_j \notin R_i$ , it is not certain that the region  $R_j$  generated from  $D_j$  is adjacent to  $R_i$  — we could miss some intermediate region. We modify our traversal algorithm as follows. The working set contains pairs  $(R, D')$  where

<sup>3</sup> There exists another flavor of PLP with parameters in the right-hand sides of the constraints.

---

**Algorithm 1** Concurrent push on the shared region structure.
 

---

```

procedure PUSH_REGION( $R$ )
  atomic ( $i \leftarrow n_{\text{fill}}; n_{\text{fill}} \leftarrow n_{\text{fill}} + 1$ )
   $regions[i] \leftarrow R$ 
  while  $n_{\text{ready}} < i$  do
    possibly use a condition variable in-
    stead of spinning
  end while  $\triangleright n_{\text{ready}} = i$ 
  atomic  $n_{\text{ready}} \leftarrow i + 1$ 
end procedure

```

---

**Algorithm 2** Task for parallel linear programming solver.
 

---

$push\_tasks$  adds new tasks to be processed (different under TBB and OpenMP).  
 $test\_and\_insert(T, x)$  checks whether  $x$  already belongs to the hash table  $T$ , in which case it returns **true**; otherwise it adds it and returns **false**. This operation is atomic.

```

procedure PROCESS_TASK( $(R_{\text{from}}, D)$ )
   $R_{\text{cov}} \leftarrow is\_covered(D, regions)$ 
  if  $R_{\text{cov}} == \text{none}$  then
     $basis \leftarrow float\_lp(A, B, C(D))$ 
    if  $\neg test\_and\_insert(bases, basis)$ 
    then
       $X^* \leftarrow exact\_point(basis)$ 
       $o \leftarrow exact\_objective(basis)$ 
      if  $\neg(X^* \geq 0 \wedge o \leq 0)$  then
         $(basis, X^*) \leftarrow exact\_lp(A, B, C(D))$ 
      end if
       $S \leftarrow sign\_conditions(basis)$ 
       $R \leftarrow eliminate\_redundancy(S)$ 
      for each constraint  $i$  in  $R$  do
         $D_{\text{next}} \leftarrow compute\_next(R, i)$ 
         $push\_tasks(D_{\text{next}})$ 
      end for
       $push\_region(R, X^*)$ 
       $R_{\text{cov}} \leftarrow R$ 
    end if
  end if
  if  $\neg are\_adjacent(R_{\text{from}}, R_{\text{cov}})$  then
     $D' \leftarrow midpoint(R_{\text{from}}, R_{\text{cov}}, D)$ 
     $W \leftarrow W \cup \{(R_{\text{from}}, D')\}$ 
  end if
end procedure

procedure  $is\_covered(D, regions)$ 
  for  $i \in 0 \dots n_{\text{ready}} - 1$  do  $\triangleright n_{\text{ready}}$  to
    be read at every loop iteration
     $(R, X^*) \leftarrow regions[i]$ 
    if  $D$  covered by  $R$  then
      return( $R$ )
    end if
  end for
  return(none)
end procedure

```

---

$R$  is a region and  $D' \notin R$  a vector (there is a special value **none** for  $R$ ). The region  $R'$  corresponding to  $D'$  is computed. If  $R$  and  $R'$  are not adjacent, then a vector  $D''$  in between  $R$  and  $R'$  is computed, and  $(R, D'')$  added to the working set. This ensures that we obtain a quasi-partition in the end. Additionally, we obtain a spanning tree of the region graph, with edges from  $R$  to  $R'$ .

The last difficulty is degeneracy. We have so far assumed that each optimization direction corresponds to exactly one basis. In general this is not the case, and the interiors of the optimality regions may overlap. This hinders performance. The final result is no longer a quasi-partition, but instead just a covering of the parameter space—enough for projection, convex hull etc. being correct.

### 3 Parallel parametric linear programming

Our algorithms are designed in a task-based execution model. The sequential algorithm executes tasks taken from a working set, which can themselves spawn new tasks. In addition, it maintains the set  $regions$  of regions already seen, used:

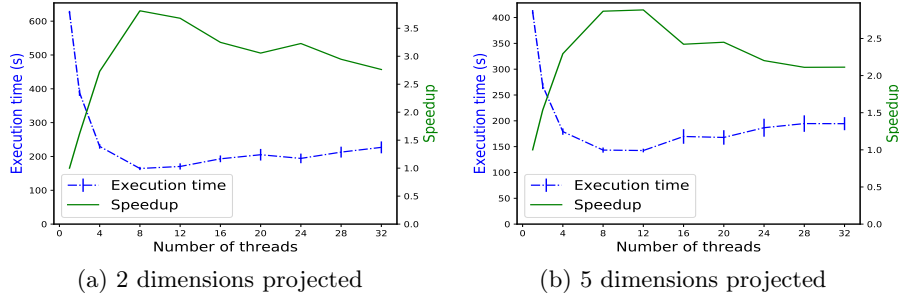


Fig. 1: 9 constraints, no redundant ones, 16 variables, 2–36 regions, OpenMP.

i) for checking if a vector  $D$  belongs to a region already covered (*is\_covered*);  
 ii) for checking adjacency of regions; iii) for adding new regions found. Therefore, in a parallel task model, this algorithm is straightforwardly parallel. The regions are inserted into a concurrent array. We investigated two task scheduling strategies. A *static* approach starts all the available tasks, waits for them to complete and collects all the new tasks  $(R, D)$  into the working set, until no new task is created and the working set is empty. A *dynamic* approach allows adding new tasks to the working set dynamically and runs the tasks until that set is empty.

The number of tasks running to completion (not aborted early due to a test) is the same as the number of generated regions. The *is\_covered* $(D, regions)$  loop can be easily parallelized as well. We opted against it as it would introduce a difficult-to-tune second level of parallelism.


We implemented these algorithms using Intel’s Thread Building Blocks (TBB [8]) and OpenMP tasks [1], both providing a task-based parallelism model with different features.

The dynamic task queue can be implemented using TBB’s `tbb::parallel_do`, which dynamically schedules tasks from the working set on a number of threads. The static scheduling approach can simply be implemented by a task synchronization barrier (such as OpenMP’s barrier).

That first implementation of the dynamic task scheduling approach was slow. The working set often contained tasks such that the regions generated from them were the same, leading to redundant computations. The workaround was to add a hash table storing the set of bases (each being identified by the ordered set of its basic variables) that have been or are currently being processed. A task will abort after solving the floating-point linear program if it finds that its basis is already in the table.


## 4 Performance evaluation

We implemented our parallel algorithms in C++, with three alternate schemes selectable at compile-time: no parallelism, OpenMP parallelism or TBB.



(a) 2 dimensions projected, 4 redundant constraints (b) 5 dimensions projected, 4 redundant constraints

Fig. 2: 24 constraints, 10 variables, 8–764 regions, OpenMP.



(a) OpenMP on Paranoia

(b) TBB on Pressembois

Fig. 3: 120 constraints, 50 variables, 1 dimension projected, 3459–3718 regions.

All benchmarks were run on the Paranoia cluster of Grid’5000 [2] and on a server called Pressembois. Paranoia has 8 nodes, each with 2 Intel<sup>®</sup> Xeon<sup>®</sup> E5-2660v2 CPUs (10 cores, 20 threads/CPU) and 128 GiB of RAM. Code was compiled using GCC 6.3.1 and OpenMP 4.5 (201511). The nodes run Linux Debian Stretch with a 4.9.0 kernel. Pressembois has 2 Intel Xeon Gold 6138 CPU (20 cores/CPU, 40 threads/CPU) and 192 GiB of RAM. It runs a 4.9 Linux kernel, and we used GCC 6.3. Every experiment was run 10 times. The plots presented in this section provide the average and standard deviation. Paranoia was used for the OpenMP experiments, whereas Pressembois was used for TBB.

We evaluated our parallel parametric linear programming implementation by using it to project polyhedra, a very fundamental operation. We used a set of typical polyhedra, with different characteristics: numbers of dimensions, of dimensions to be projected and of constraints, sparsity. Here we present a subset of these benchmarks, each comprising 50 to 100 polyhedra.

On problems that have only few regions, not enough parallelism can be extracted to exploit all the cores of the machine. For instance, Figure 1 presents two experiments on 2 to 36 regions using the OpenMP version. It gives an acceptable speed-up on a few cores (up to 10), then the computation does not generate enough tasks to keep the additional cores busy. As expected, when the solution

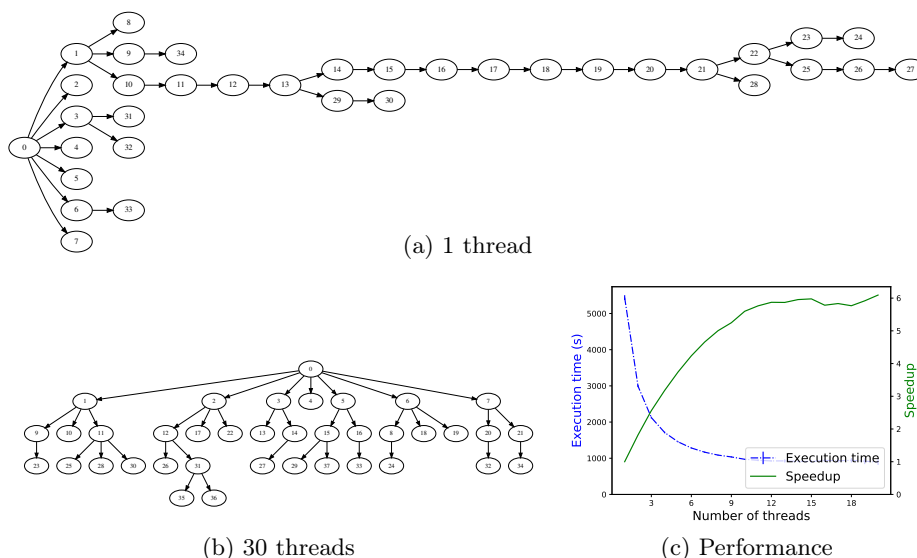


Fig. 4: Generation graph of the regions from one polyhedron, computed with 1 and 30 threads. The region graphs, depending on overlaps etc., are different; the numbers in both trees have no relationship.

has many regions, computation scales better. Figure 2 presents the performance obtained on polyhedra made of 24 constraints, involving 8 to 764 regions, using the OpenMP version. The speed-up is sublinear, especially beyond 20 cores.

On larger polyhedra, with 120 constraints and 50 variables, the speedup is close to linear with both OpenMP and TBB (Fig. 3). The parallelism extracted from the computation is illustrated by Fig. 4, on a polyhedron involving 29 constraints and 16 variables. Figure 4b shows the number of parallel tasks.

## References

1. OpenMP Application Programming Interface, 4.5 edn. (2015)
2. Cappello, F., al: Grid’5000: A large scale and highly reconfigurable grid experimental testbed. In: International Workshop on Grid Computing. IEEE/ACM (2005)
3. Cousot, P., Halbwachs, N.: Automatic discovery of linear restraints among variables of a program. In: POPL. pp. 84–96. ACM Press (1978)
4. Jones, C., al: On polyhedral projections and parametric programming. *J. Optimization Theory and Applications* **138**(2), 207–220 (2008)
5. Jones, C.N., Kerrigan, E.C., Maciejowski, J.M.: Lexicographic perturbation for multiparametric linear programming with applications to control. *Automatica* (2007)
6. Maréchal, A., Périn, M.: Efficient elimination of redundancies in polyhedra by ray-tracing. In: VMCAI. LNCS, vol. 10145, pp. 367–385. Springer (2017)
7. Maréchal, A., Monniaux, D., Périn, M.: Scalable minimizing-operators on polyhedra via parametric linear programming. In: SAS. Springer (2017)
8. Reinders, J.: Intel threading building blocks: outfitting C++ for multi-core processor parallelism. " O’Reilly Media, Inc." (2007)