Parallel Constraint-Driven Inductive Logic Programming

Andrew Cropper,[1] **Oghenejokpeme Orhobor,** [2] **Cristian Dinu,** [1] **Rolf Morel** [1]

[1] University of Oxford
[2] University of Cambridge
andrew.cropper@cs.ox.ac.uk, oo288@cam.ac.uk, cristian.dinu@hertford.ox.ac.uk, rolf.morel@cs.ox.ac.uk

Abstract

Multi-core machines are ubiquitous. However, most inductive logic programming (ILP) approaches use only a single core, which severely limits their scalability. To address this limitation, we introduce parallel techniques based on *constraint-driven* ILP where the goal is to accumulate constraints to restrict the hypothesis space. Our experiments on two domains (program synthesis and inductive general game playing) show that (i) parallelisation can substantially reduce learning times, and (ii) worker communication (i.e. sharing constraints) is important for good performance.

1 Introduction

Inductive logic programming (ILP) (Muggleton 1991) is a form of machine learning. Given positive and negative examples of a predicate and background knowledge (BK), the ILP problem is to find a set of logical rules (a *hypothesis*), which, with the BK, entails the positive examples and none of the negative examples. Key features of ILP include its ability to (i) learn from small amounts of data, (ii) support relational data, and (iii) induce human-readable hypotheses (Cropper et al. 2021).

Although powerful, ILP approaches often struggle with scalability: efficiently searching a large hypothesis space (the set of all hypotheses) for a *solution* (a hypothesis that correctly generalises the examples). Parallelisation is one approach to improving scalability. However, although multi-core machines are ubiquitous, most ILP approaches are single-core learners (Muggleton 1995; Srinivasan 2001; Cropper and Muggleton 2016; Law, Russo, and Broda 2014).

To overcome this limitation, we introduce parallel techniques based on *constraint-driven* ILP where the goal is to accumulate constraints to restrict the hypothesis space. In particular, we build on the *learning from failures* (LFF) (Cropper and Morel 2021a,b) approach, which supports predicate invention and learning optimal and recursive programs. A LFF learner works by repeatedly generating and testing hypotheses on training examples. If a hypothesis fails to correctly generalise the examples, the learner deduces constraints to *explain* the failure, which it then uses to rule out other hypotheses and thus restrict the hypothesis space. The process repeats until a solution is found.

We introduce two general constraint-driven parallel ILP approaches based on parallel *conflict-driven clause learning* (CDCL) SAT techniques (Martins, Manquinho, and Lynce 2012), namely *portfolio* and *divide-and-conquer* approaches. In our portfolio approach, parallel LFF learners compete by searching the same hypothesis space using different heuristics. In our divide-and-conquer approach, we divide the hypothesis space into disjoint subspaces which we assign to parallel LFF learners who each search them. Figure 1 illustrates these strategies. We also allow learners to exchange learned constraints, similar to how parallel SAT techniques share clauses.

Overall, our contributions are:

- We introduce parallel ILP approaches inspired by *portfolio* and *divide-and-conquer* approaches used by CDCL SAT solvers.
- We implement the techniques to parallelise the LFF implementation Popper.
- We experimentally show on two domains (program synthesis and inductive general game playing) that (i) our parallel methods can lead to linear speedups with up to four processors in general, (ii) our parallel methods can lead to super-linear speedup in some cases, and (iii) that communication (i.e. sharing constraints) is important for good performance.

2 Related Work

2.1 Sequential ILP

Many ILP systems, such as Progol (Muggleton 1995) and Aleph (Srinivasan 2001), often struggle with large numbers of examples because of their sequential set covering approach. A notable exception is QuickFOIL (Zeng, Patel, and Page 2014) which builds on FOIL (Quinlan 1990) by introducing (i) a new scoring function for clauses, and (ii) a highly efficient relational database implementation. The authors show that their approach can scale to datasets with millions of background facts and hundreds of thousands of examples. However, because it builds on FOIL, QuickFOIL inherits its limitations, including (i) difficulty learning recursive programs, (ii) no support for predicate invention, and (iii) no guarantees about the optimality of solutions. These limitations apply to almost all classical ILP systems (Cropper et al. 2021).
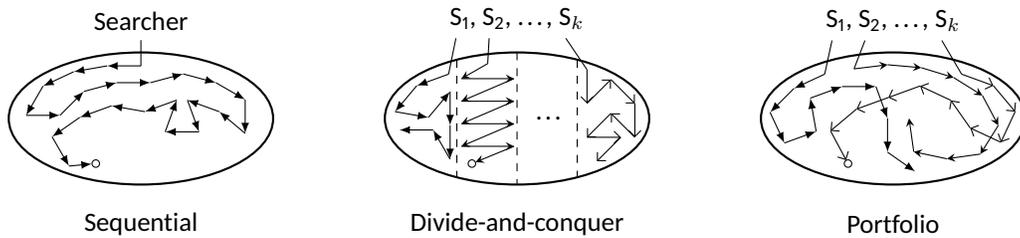
Figure 1: One sequential and two parallel search strategies for finding a solution in the same hypothesis space.

Modern meta-level ILP approaches (Corapi, Russo, and Lupu 2011; Cropper and Muggleton 2016; Law, Russo, and Broda 2014; Kaminski, Eiter, and Inoue 2018; Evans et al. 2021; Cropper and Morel 2021a) can often learn recursive and optimal programs and perform predicate invention. Although there is no standard definition for *meta-level ILP* most approaches encode the ILP problem as a meta-level logic program, i.e. a program that reasons about programs. These approaches delegate the search for a hypothesis to an off-the-shelf solver, such as an answer set programming (ASP) solver (Gebser et al. 2014), after which the meta-level solution is translated back to a standard solution for the ILP task. However, modern meta-level ILP approaches often struggle in terms of scalability. For instance, ILASP (Law, Russo, and Broda 2014) struggles to learn rules with lots of literals because it precomputes every possible rule that may appear in a program, which is often infeasible. The same issue prevents HEXMIL (Kaminski, Eiter, and Inoue 2018) from scaling to large problems. As far as we are aware, there is no parallel meta-level ILP system.

2.2 Parallel ILP

Fonseca et al. (2009) survey parallel ILP techniques. They divide approaches into three categories: *search*, *data*, and *evaluation*. Search approaches parallelise the search of the hypothesis space. Our approaches are in this category. Data approaches divide the training examples amongst the workers which learn solutions for them in parallel. Evaluation approaches evaluate candidates rules in parallel.

We discuss three notable parallel approaches. Dehaspe and De Raedt (1995) parallelise the search for a hypothesis in the ILP system Claudien. Claudien works by maintaining a priority queue of potential clauses to add to a hypothesis. If a clause is too general, it is removed from the queue and its specialisations are added. In the parallel approach, parallel workers process clauses in the queue. The authors experimentally show that the parallelisation speed-up is roughly proportional to the number of workers.

Wang and Skillicorn (2000) parallelise Progol by allocating a subset of the positive examples (and all of the negative examples) to each worker. Each worker applies the standard Progol sequential algorithm to find a good clause for its subset of the positive examples, which it then communicates to the other workers, who may then incorporate the clause into their local theory. The authors show linear speedups with four and six processors.

Srinivasan, Faruquie, and Joshi (2012) use the MapReduce paradigm to parallelise the scoring step of Aleph, where a clause is generated and its score is calculated as a function of the examples. Their results are generally positive, especially when the number of examples is large. However, although the authors rightly claim that their approach is not specific to Aleph, it is specific to the classic set covering approach.

These three approaches are all limited because they inherit the same limitations as their sequential counterparts. By contrast, a key contribution of this paper is the introduction of parallel techniques that can learn optimal and recursive programs and perform predicate invention.

More recently, Katzouris, Artikis, and Paliouras (2017) introduce p-OLED, a parallel version of OLED (Katzouris, Artikis, and Paliouras 2016), which learns event definitions in the form of even calculus theories. In their parallel approach, a clause is evaluated in parallel on sub-streams of the input stream and its independent scores are combined. Their evaluations show that their approach can reduce training times and, in some cases, is capable of super-linear speed-ups. In contrast to p-OLED, our approaches learn general definite programs, including programs with recursion and predicate invention.

2.3 Parallel SAT

Our parallel ILP approaches are based on parallel conflict-driven-clause-learning SAT techniques (Martins, Manquinho, and Lynce 2012), of which there are two main approaches. *Divide-and-conquer* approaches divide the search space into sub-spaces which are allocated to sequential workers. Workers co-operate through the exchange of learnt conflicts. Portfolio approaches (Hamadi, Jabbour, and Sais 2009) allow multiple sequential workers to compete on the same search space by employing different search heuristics, such as different restart policies, branching heuristics, random seeds, etc. Workers also co-operate through the exchange of learnt conflicts. This exchange can be done through message passing (Schubert, Lewis, and Becker 2009), which is necessary for distributed approaches, or through shared memory (Hamadi, Jabbour, and Sais 2009). In this paper, we transfer these high-level ideas to the area of constraint-driven ILP. In our approach, the search space is the hypothesis space and workers can co-operate through the exchange of learned constraints, i.e. clauses that describe conflicts.

3 Problem Setting

We now define the LFF problem, on which the approaches in Section 4 are based. We assume familiarity with logic programming (Lloyd 2012).

The key idea of LFF is to use *hypothesis constraints* to restrict the hypothesis space. Let $\mathcal{L}$ be a language that defines hypotheses, i.e. a meta-language. For instance, consider a meta-language formed of two literals, *h_lit/4* and *b_lit/4*, which represent *head* and *body* literals respectively. With this language, we can denote the clause *last(A,B) ← tail(A,C), head(C,B)* as the set of literals {*h_lit(0,last,2,(0,1)), b_lit(0,tail,2,(0,2)), b_lit(0,head,2,(2,1))*}. The first argument of each literal is the clause index, the second is the predicate symbol, the third is the arity, and the fourth is the literal variables, where *0* represents *A*, *1* represents *B*, etc.

A *hypothesis constraint* is a constraint expressed in $\mathcal{L}$. Let C be a set of hypothesis constraints written in a language $\mathcal{L}$. A set of definite clauses H is *consistent* with C if, when written in $\mathcal{L}$, H does not violate any constraint in C. For instance, the constraint ← *h_lit(0,last,2,(0,1)), b_lit(0,last,2,(1,0))* would be violated by the definite clause *last(A,B) ← last(B,A)*. Let $\mathcal{H}$ be a hypothesis space. We denote as $\mathcal{H}_C$ the subset of $\mathcal{H}$ which do not violate any constraint in C.

We define the LFF problem:

Definition 1 (**LFF input**). The *LFF* input is a tuple $(E^+, E^-, B, \mathcal{H}, C)$ where E^+ and E^- are sets of ground atoms denoting positive and negative examples respectively; B is a Horn program denoting background knowledge; $\mathcal{H}$ is a hypothesis space, and C is a set of hypothesis constraints.

We define a LFF solution:

Definition 2 (**LFF solution**). Given an input tuple $(E^+, E^-, B, \mathcal{H}, C)$, a hypothesis $H \in \mathcal{H}_C$ is a *solution* when H is *complete* ($\forall e \in E^+, \ B \cup H \models e$) and *consistent* ($\forall e \in E^-, \ B \cup H \not\models e$).

If a hypothesis is not a solution then it is a *failure* or a *failed hypothesis*. A hypothesis is *incomplete* when $\exists e \in E^+, \ H \cup B \not\models e$. A hypothesis is *inconsistent* when $\exists e \in E^-, \ H \cup B \models e$. A hypothesis is *totally incomplete* when $\forall e \in E^+, \ H \cup B \not\models e$.

Let $cost : \mathcal{H} \mapsto R$ be an arbitrary cost function. We define an *optimal* solution:

Definition 3 (**Optimal solution**). Given an input tuple $(E^+, E^-, B, \mathcal{H}, C)$, a hypothesis $H \in \mathcal{H}_C$ is *optimal* when (i) H is a solution, and (ii) $\forall H' \in \mathcal{H}_C$, where H' is a solution, $cost(H) \leq cost(H')$.

In this paper, we define the *cost(H)* to be the total number of literals in the logic program H.

Hypothesis Constraints Cropper and Morel (2021a,b) introduce hypothesis constraints based on subsumption (Plotkin 1971). A clause C_1 *subsumes* a clause C_2 ($C_1 \preceq C_2$) if and only if there exists a substitution θ such that $C_1\theta \subseteq C_2$. A clausal theory T_1 subsumes a clausal theory T_2 ($T_1 \preceq T_2$) if and only if $\forall C_2 \in T_2, \exists C_1 \in T_1$ such that C_1 subsumes C_2. A clausal theory T_1 is a *specialisation* of

a clausal theory T_2 if and only if $T_2 \preceq T_1$. A clausal theory T_1 is a *generalisation* of a clausal theory T_2 if and only if $T_1 \preceq T_2$. If a hypothesis H is incomplete, a *specialisation* constraint prunes specialisations of H, as they are guaranteed to also be incomplete. If a hypothesis H is inconsistent, a *generalisation* constraint prunes generalisations of H, as they are guaranteed to be inconsistent as well. If a hypothesis H is totally incomplete, a *redundancy* constraint prunes hypotheses that contain a specialisation of H as a subset.

4 Parallel Algorithms

We now describe our parallel ILP approaches. We first describe the sequential ILP system Popper, which we parallelise.

4.1 Popper

Algorithm 1 shows the high-level Popper algorithm, which solves the LFF problem (Definition 1). Popper takes as input positive (pos) and negative (neg) examples, background knowledge (bk), and a maximum hypothesis size (max_size). Popper uses a *generate*, *test*, and *constrain* loop. Popper starts with a base *generator* ASP program whose models correspond to hypotheses (definite programs). The idea is to augment this generator program with constraints to eliminate models and thus restrict the hypothesis space. The constraints are initially empty (line 3).

In the generate stage (line 5), Popper uses Clingo (Gebser et al. 2014), an ASP system, to search for a model of the generator program with exactly m literals which Popper then converts to a hypothesis (a definite program).

In the test stage (line 9), Popper tests a hypothesis on the given training examples. If a hypothesis fails, i.e. is *incomplete* or *inconsistent*, then, in the constrain stage (line 12), Popper learns hypothesis constraints (described as ASP constraints) from the failure. Popper adds the constraints to the generator program to prune models and constrain subsequent hypothesis generation. For instance, if a hypothesis is incomplete, i.e. does not entail all the positive examples, Popper builds a specialisation constraint to prune hypotheses that are logically more specific.

To find an optimal solution, Popper progressively increases the number of literals allowed in a hypothesis when the hypothesis space is empty at a certain size (e.g. when the generator program together with the learned constraints has no more models) (line 6). This loop repeats until either (i) Popper finds an optimal solution, or (ii) there are no more hypotheses to test. Popper is guaranteed to find the optimal solution when every hypothesis is guaranteed to terminate, such as when the hypothesis space only contains Datalog programs.

4.2 Parallel Solving

The simplest way to parallelise Popper is to parallelise the search for a model in the generate stage using the parallel capabilities of Clingo (Gebser, Kaufmann, and Schaub 2012). Clingo incorporates a SAT solver which supports *parallel search* using shared memory multi-threading (cf. Section

```
Algorithm 1: Popper
1  def popper(pos, neg, bk, max_size):
2    m = 1
3    cons = {}
4    while m ≤ max_size:
5      h = generate(cons, m)
6      if h == UNSAT:
7        m += 1
8      else:
9        outcome = test(pos, neg, bk, h)
10       if outcome == (COMPLETE, CONSISTENT)
11         return h
12       cons += constrain(h, outcome)
13   return UNSAT
```

```
Algorithm 2: Portfolio
1  def port(pos, neg, bk, max_size, q_cons, com):
2    m = 1
3    cons = {}
4    while m ≤ max_size:
5      h = generate(cons, m)
6      if h == UNSAT:
7        m += 1
8      else:
9        outcome = test(pos, neg, bk, h)
10       if outcome == (COMPLETE, CONSISTENT)
11         return h
12       cons' = constrain(h, outcome)
13       cons += cons'
14       if com:
15         q_cons.put(cons')
16         cons += q_cons.get()
17   return UNSAT
```

2.3). Learned conflict clauses (known as *nogoods*) are exchanged among worker threads according to various heuristics. To implement this approach, we simply enable Clingo's parallel mode (using the flag '–parallel-mode k') which runs k workers in a portfolio configuration. In our experiments, we call this approach ParSearch.

4.3 Portfolio

Similar to parallel portfolio SAT approaches, our parallel portfolio ILP approach involves multiple workers searching the same hypothesis space using different strategies. Algorithm 2 shows the main portfolio worker algorithm, which we call Portfolio and which is almost identical to Popper (Algorithm 1). A master controls the workers. The master first creates an empty message queue to allow workers to communicate. This queue is a *many-to-many* queue that allows each worker to receive a copy of any message put on it. The master then spawns k workers who each search the same hypothesis space.

In addition to the standard Popper inputs, a Portfolio worker receives as input a message queue for communicating constraints (q_cons) and a flag to denote whether to send constraints to other workers. Each worker calls the generate step (i.e. Clingo) with different search heuristics so that they find models (and thus hypotheses) in a different order to the other workers. In this paper, we use the simple approach of calling Clingo with the arguments *--rand-freq=p* and *--seed=s*. The *--rand-freq* flag tells Clingo to perform random (rather than heuristic) decisions with probability p. The value s is a seed. We set $p = 0.01$ and s to be the workerid. A direction for future work is to determine how to choose a suitable Clingo heuristic. To be clear, each Portfolio worker has its own single-threaded ASP solver. As with Popper, Portfolio is guaranteed to find the optimal solution when every hypothesis is guaranteed to terminate.

Portfolio$_{com}$ The Portfolio approach, by default, does not permit communication between workers (i.e. the communication flag is false by default). If the communication option (com) is true, then workers exchange constraints with other workers. We call this communication-enabled version of the algorithm Portfolio$_{com}$. When one Portfolio$_{com}$ worker finds an incomplete or inconsistent hypothesis, it builds the constraints and sends them to the other workers (line 15) and

also receives them (line 16). As Popper is essentially learning nogoods by testing hypotheses, our sharing of learned constraints can be seen as exchanging *externally learnt* nogoods between multiple solvers.

4.4 Divide and Conquer

Similar to parallel D&C SAT approaches, our parallel D&C ILP approach involves multiple workers searching disjoint hypothesises spaces. In this paper, we take the simple approach of splitting the hypothesis space by the hypothesis size. A direction for future work is to study alternative methods to divide the hypothesis space.

Algorithm 3 shows the D&C worker algorithm, which we call D&C. As with Portfolio, a master controls the workers. In addition to creating an empty message queue to share constraints between workers (q_cons), the master creates a second message queue (q_size) to maintain hypothesis sizes to be searched. The master populates the size queue with all possible hypothesis sizes in increasing order. The master then spawns k workers.

In addition to the standard Popper inputs, a D&C worker receives as input a message queue for communicating constraints (q_cons), a message queue for receiving hypothesis sizes to explore (q_sizes), and a flag to denote whether to send constraints to other workers. A worker pops the smallest size m off the size queue and then searches for a solution with exactly m literals. If there is no solution, the worker loops again and pops another size off the queue. Since each worker only considers solutions with exactly m literals, they all consider disjoint regions of the hypothesis space. If a worker finds a solution, they inform the master. If a solution of size m is found, the master waits until all workers searching for hypotheses of size $k < m$ have finished before returning the solution to guarantee that D&C returns an optimal solution.

D&C$_{com}$ The D&C approach, by default, does not permit communication between workers (i.e. the communication is defaulted to false). If the communication option (com) is true, then workers exchange constraints with other work-

```
Algorithm 3: D&C
1 def dac(pos, neg, bk, q_size, q_cons, com):
2   cons = {}
3   while |q_size| > 1:
4     m = q_size.get()
5     while true:
6       h = generate(cons, m)
7       if h == UNSAT:
8         break
9       outcome = test(pos, neg, bk, h)
10      if outcome == (COMPLETE, CONSISTENT)
11        return h
12      cons' = constrain(h, outcome)
13      cons += cons'
14      if com:
15        q_cons.put(cons')
16        cons += q_cons.get()
17   return UNSAT
```
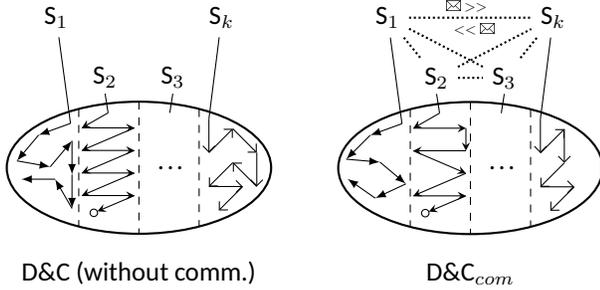

Figure 2: Divide-and-conquer strategy with and without communication. With communication enabled, learned constraints are passed as messages between workers.

ers. We call this communication-enabled version of the algorithm D&C$_{com}$. Figure 2 illustrates the difference between D&C and D&C$_{com}$.

4.5 Implementation

We implemented the systems in Python 3 using the multi-processing module. Each process is an instance of Popper which maintains its own ASP solver. Rather than use shared memory, we use a message queue to allow workers to communicate.

5 Experiments

We claim that our parallelisation approaches can reduce learning times and thus improve scalability. Our experiments[1] therefore aim to answer the question:

Q1 Can our parallel approaches reduce learning times?

To answer this question, we compare the performance of our approaches when given progressively more workers (CPUs).

---

[1]All the implementation code and experimental data will be made open-source and freely available after the paper has gone through review.

```
f(A,B):- empty(A),empty(B).
f(A,B):- head(A,D),odd(D),tail(A,C),f(C,B).
f(A,B):- tail(A,C),head(A,E),even(E),
         f(C,D),prepend(E,D,B).
```

Figure 3: An example solution described as Prolog program for the *filter* task.

Our parallel approaches allow workers to exchange learned constraints. To evaluate the impact of communication, our experiments aim to answer the question:

Q2 Can communication (sharing constraints) reduce learning times?

Settings. We use Clingo version 5.5.0, SWI-Prolog 7.6.3 and Python 3.9.6. All experiments were performed on a server with an AMD Opteron™ Processor with 32 cores, 64 threads, and 256GB of RAM. We give all the approaches identical inputs in all the experiments. We enforce a time-out of five minutes per task. We repeat each experiment five times and measure mean learning time and standard deviation.

5.1 Domains

We consider two domains: *program synthesis* and *inductive general game playing*.

Program synthesis Program synthesis has long been considered a difficult problem in ILP (Muggleton et al. 2012). Indeed, most ILP systems cannot learn recursive programs. Cropper and Morel (2021a) showed that Popper can learn recursive programs with higher accuracies and in less time than other ILP systems. In this experiment, we try to answer the experimental questions using a similar program synthesis dataset.

We use four challenging synthesis tasks: (*find dupl*) find a duplicate element in a list, (*sorted*) determine whether a list is sorted, (*dropk*) drop the first k elements in a list, and (*filter*) remove all odd elements from a list. Figure 3 shows an example solution for the *filter* task. We provide as background knowledge the dyadic relations *head*, *tail*, *element*, *increment*, *decrement*, *geq* and the monadic relations *empty*, *zero*, *one*, *even*, and *odd*. We also include the triadic relation *prepend* in the *filter* experiment.

For each task, we generate 10 positive and 10 negative training examples. Each list is randomly generated and has a maximum length of 50. We sample the list elements uniformly at random from the set $\{1, 2, \ldots, 100\}$.

IGGP The general game playing (GGP) framework (Genesereth and Björnsson 2013) is a system for evaluating an agent's general intelligence across a wide range of tasks. In the GGP competition, agents are tested on games they have never seen before. In each round, the agents are given the rules of a new game. The rules are described symbolically as a logic program. The agents are given a few seconds to process the rules of the game, then they start playing, thus producing game traces. The

```
nextscore(A,B,C):- score(A,B,C),does(A,D,F),
        does(A,E,F),different(D,E).
nextscore(A,B,C):- different(B,E),does(A,E,G),
        score(A,F,C),does(A,F,D),beats(G,D).
nextscore(A,B,C):- score(A,B,F),beats(E,G),
        does(A,D,G),succ(F,C),
        different(B,D),does(A,B,E).
```

Figure 4: An example solution described as Prolog program for the *rps* task.

winner of the competition is the agent who gets the best total score over all the games. In this experiment, we use the IGGP dataset (Cropper, Evans, and Law 2020) which inverts the GGP task: an ILP system is given game traces and the task is to learn a set of rules (a logic program) that could have produced these traces. We focus on two IGGP games: *minimal decay* and *rock, paper, scissors* (rps) and on learning the *next* relations, which is the most challenging one to learn (Cropper, Evans, and Law 2020). Figure 4 shows an example solution for the *rps* task.

5.2   Results

Figure 5 shows the experimental results. Popper is the single-core baseline.

ParSearch.   The results show no major benefit from ParSearch, i.e. running Popper with parallel Clingo enabled. For instance, in the *filter* experiment, the difference in running time between Popper and ParSearch is less than 5%. This result may surprise the reader, especially as parallel Clingo has been shown to outperform sequential Clingo (Gebser, Kaufmann, and Schaub 2012). The reason is that the bottleneck in Popper is rarely generating a hypothesis, i.e. searching for a syntactically valid program. Instead, the bottleneck is mostly the sheer number of hypotheses to consider. There is, therefore, little benefit from parallelising *only* the generate part.

Portfolio.   The results are mixed for the Portfolio approach without communication. In two tasks it performs worse than Popper (*dropk* and *sorted*), in one better (*minimal decay*), and about the same in the rest. This result is expected. The only way that Portfolio can outperform Popper is when a worker happens to find a solution quicker than Popper because of its different search heuristics. However, since Popper performs iterative deepening search on the hypothesis sizes, whereby it proves that there is no solution at a certain size before going to the next size, both Popper and Portfolio should roughly take the same amount of time to search the space of programs smaller than the solution. It is only when searching the part of the space at the size of the solution that Portfolio has the opportunity to outperform Popper, hence the modest improvements.

Portfolio$_{com}$.   The results are very strong for the Portfolio$_{com}$ approach. In all cases, Portfolio$_{com}$ outperforms Popper and a paired t-test confirms the significance at the $p < 0.01$ level. For instance, in the *minimal decay*

experiment, given one worker, Portfolio$_{com}$ takes about the same time (around 250s) as Popper, which is to be expected as there is no parallelisation with only one worker. Given more workers, Portfolio$_{com}$ starts to outperform Popper. With two workers, the learning time of Portfolio$_{com}$ is almost halved to 136s. With four workers, the learning time of Portfolio$_{com}$ is halved again to 60s. With eight workers, the learning time of Portfolio$_{com}$ is almost halved again to around 35s. Similar reductions are demonstrated in all the tasks. This result strongly suggests that the answer to **Q1** is yes: our parallelisation approaches can significantly reduce learning times and that the speed-up is roughly proportional to the number of workers. In all cases, Portfolio$_{com}$ outperforms Portfolio and a paired t-test confirms the significance at the $p < 0.01$ level. This result clearly demonstrates that communication between workers is important for good learning performance (**Q2**).

D&C.   Figure 5 excludes the results of D&C (D&C without communication) results. We have excluded the results for D&C because it struggles to find solutions for any tasks in the given time limit. This result may surprise the reader, as they may think that D&C would in the worst-case simulate Popper. However, the key omission of D&C is the inability to learn constraints from failed hypotheses from multiple hypothesis sizes. A key reason why Popper can efficiently find solutions is that it first considers smaller hypotheses before larger ones. By learning constraints from small failed hypotheses, Popper can prune large parts of the hypothesis space. So when Popper increases the hypothesis size bound, the space for the next size is greatly reduced from constraints learnt when learning for smaller sizes. However, in the D&C approach, there is no transfer of knowledge between the hypothesis sizes. By contrast, D&C$_{com}$ performs reasonably well, outperforming Popper on four of the six tasks. The most impressive results are in the *minimal decay* task. With one worker, D&C$_{com}$ takes around 250s. With two workers, the running time of D&C$_{com}$ is reduced to one fifth (45s). Given eight workers, D&C$_{com}$ takes around 1s to find a solution. This result again clearly demonstrates that communication between workers is important for good learning performance (**Q2**).

6   Conclusions and Limitations

To improve the ILP scalability problem, we have introduced parallel ILP approaches inspired by parallel SAT approaches, namely *portfolio* and *divide-and-conquer* approaches. Our implementations parallelise Popper, a state-of-the-art constraint-driven ILP system. Our experiments on two domains (program synthesis and inductive general game playing) show that (i) our parallel methods can lead to linear speedups with up to four processors in general, (ii) our parallel methods can lead to super-linear speedup in some cases, and (iii) that communication (i.e. sharing constraints) is important for good performance. As far as we are aware, this work is the first to clearly demonstrate the ability to parallelise state-of-the-art ILP systems that support predicate invention and learning recursive programs.
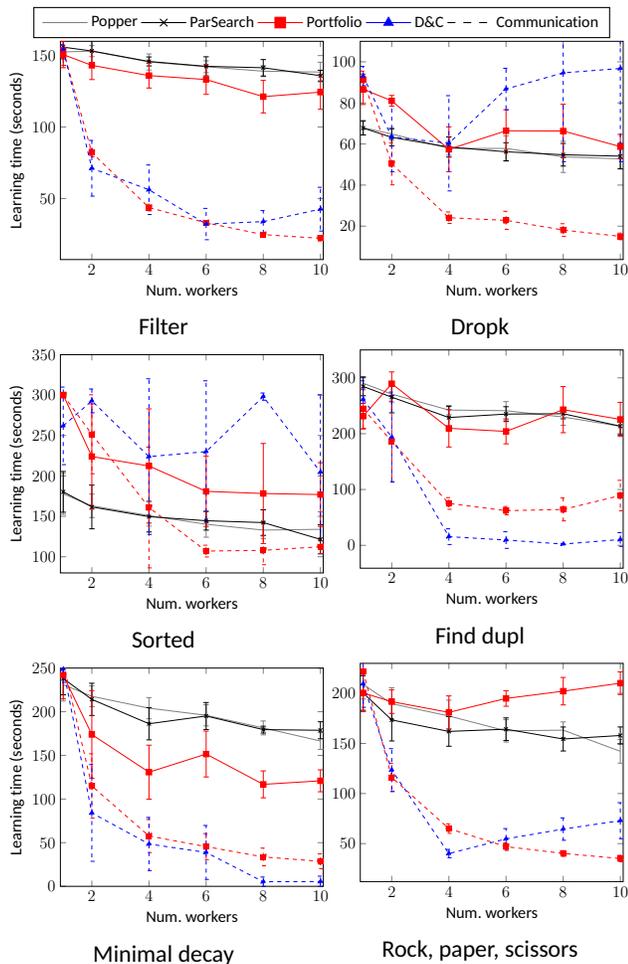
Figure 5: Experimental results.

6.1 Limitations and Future Work

There are many limitations of this work, which open much scope for future work.

Better implementation. Our implementations are based on message passing via inter-process communication. As our workers learn many, and at times large, constraints, message passing involves non-trivial overhead, not least due to copying. On a single machine, we could exploit shared memory for improved performance, e.g. due to less copying.

Better constraints. Parallel SAT solver workers tend to perform checks on the clauses that they receive. One such check ensures that a new *propositional* constraint is not subsumed by any known constraint (Hamadi, Jabbour, and Sais 2009). We could also check subsumption of our learned *first-order* constraints, though this does have non-trivial cost. SAT solvers also employ heuristics for deciding which constraints to share. For example, typically only small learned constraints are sent to other workers, based on some cut-off value (Martins, Manquinho, and Lynce 2012). Another heuristic is to only send (long) constraints to workers searching a similar part of the solution space (ibid.). We

could empirically test if parallel SAT techniques like these are beneficial to parallel constraint-driven ILP.

Distributed approaches. Our implementations support local parallelisation, i.e. on one machine. Our choice for message passaging is an advantage here, as it should generalise to distributed parallelisation, where we could potentially harness hundreds of CPUs.

Combination of D&C and Portfolio The D&C approach sometimes has one worker working on size k and other workers doing nothing, as they have finished working on other sizes. Instead of them waiting, they could start work on the same size k, each with a new search heuristic. In line with a distributed approach having many workers, we could follow the parallel SAT strategy of first splitting the search space and then applying a portfolio of workers to each subspace.

References

Corapi, D.; Russo, A.; and Lupu, E. 2011. Inductive Logic Programming in Answer Set Programming. In Muggleton, S.; Tamaddoni-Nezhad, A.; and Lisi, F. A., eds., *Inductive Logic Programming - 21st International Conference, ILP 2011, Windsor Great Park, UK, July 31 - August 3, 2011, Revised Selected Papers*, volume 7207 of *Lecture Notes in Computer Science*, 91–97. Springer.

Cropper, A.; Dumancic, S.; Evans, R.; and Muggleton, S. H. 2021. Inductive logic programming at 30. *arXiv*.

Cropper, A.; Evans, R.; and Law, M. 2020. Inductive general game playing. *Machine Learning*, 109(7): 1393–1434.

Cropper, A.; and Morel, R. 2021a. Learning programs by learning from failures. *Mach. Learn.*, 110(4): 801–856.

Cropper, A.; and Morel, R. 2021b. Predicate Invention by Learning From Failures. *CoRR*.

Cropper, A.; and Muggleton, S. H. 2016. Metagol System. https://github.com/metagol/metagol.

Dehaspe, L.; and De Raedt, L. 1995. Parallel inductive logic programming. In *The MLnet Familiarization Workshop on Statistics, Machine Learning and Knowledge Discovery in Databases*.

Evans, R.; Hernández-Orallo, J.; Welbl, J.; Kohli, P.; and Sergot, M. J. 2021. Making sense of sensory input. *Artif. Intell.*, 293: 103438.

Fonseca, N. A.; Srinivasan, A.; Silva, F. M. A.; and Camacho, R. 2009. Parallel ILP for distributed-memory architectures. *Mach. Learn.*, 74(3): 257–279.

Gebser, M.; Kaminski, R.; Kaufmann, B.; and Schaub, T. 2014. Clingo = ASP + Control: Preliminary Report. *CoRR*, abs/1405.3694.

Gebser, M.; Kaufmann, B.; and Schaub, T. 2012. Multi-threaded ASP solving with clasp. *TPLP*.

Genesereth, M. R.; and Björnsson, Y. 2013. The International General Game Playing Competition. *AI Magazine*, 34(2): 107–111.

Hamadi, Y.; Jabbour, S.; and Sais, L. 2009. ManySAT: a Parallel SAT Solver. *J. Satisf. Boolean Model. Comput.*

Kaminski, T.; Eiter, T.; and Inoue, K. 2018. Exploiting An-swer Set Programming with External Sources for Meta-Interpretive Learning. *Theory Pract. Log. Program.*, 18(3-4): 571–588.

Katzouris, N.; Artikis, A.; and Paliouras, G. 2016. Online learn-ing of event definitions. *TPLP*, 16(5-6): 817–833.

Katzouris, N.; Artikis, A.; and Paliouras, G. 2017. Parallel On-line Learning of Event Definitions. In Lachiche, N.; and Vrain, C., eds., *Inductive Logic Programming - 27th International Conference, ILP 2017, Orléans, France, September 4-6, 2017, Revised Selected Papers*, volume 10759 of *Lecture Notes in Computer Science*, 78–93. Springer.

Law, M.; Russo, A.; and Broda, K. 2014. Inductive Learning of Answer Set Programs. In *JELIA*.

Lloyd, J. W. 2012. *Foundations of logic programming*. Springer Science & Business Media.

Martins, R.; Manquinho, V. M.; and Lynce, I. 2012. An overview of parallel SAT solving. *Constraints An Int. J.*

Muggleton, S. 1991. Inductive Logic Programming. *New Gen-eration Computing*, 8(4): 295–318.

Muggleton, S. 1995. Inverse Entailment and Progol. *New Generation Comput.*, 13(3&4): 245–286.

Muggleton, S.; De Raedt, L.; Poole, D.; Bratko, I.; Flach, P. A.; Inoue, K.; and Srinivasan, A. 2012. ILP turns 20 - Biography and future challenges. *Machine Learning*, 86(1): 3–23.

Plotkin, G. 1971. *Automatic Methods of Inductive Inference*. Ph.D. thesis, Edinburgh University.

Quinlan, J. R. 1990. Learning Logical Definitions from Rela-tions. *Mach. Learn.*, 5: 239–266.

Schubert, T.; Lewis, M.; and Becker, B. 2009. PaMiraXT: Par-allel SAT Solving with Threads and Message Passing. *J. Satisf. Boolean Model. Comput.*, 6(4): 203–222.

Srinivasan, A. 2001. The ALEPH manual. *Machine Learning at the Computing Laboratory, Oxford University*.

Srinivasan, A.; Faruquie, T. A.; and Joshi, S. 2012. Data and task parallelism in ILP using MapReduce. *Mach. Learn.*, 86(1): 141–168.

Wang, Y.; and Skillicorn, D. 2000. *Parallel inductive logic in data mining*. Citeseer.

Zeng, Q.; Patel, J. M.; and Page, D. 2014. QuickFOIL: Scalable Inductive Logic Programming. *Proc. VLDB Endow.*, 8(3): 197–208.