

Application of Meta-Tree-Based Distributed Search to the Railway Scheduling Problem*

Montserrat Abril, Miguel A. Salido, Federico Barber

Dpto. Sistemas Informáticos y Computación, Universidad Politécnica de Valencia
Camino de Vera s/n, 46022, Valencia, Spain
{mabril, msalido, fbarber}@dsic.upv.es

Abstract

Many problems of theoretical and practical interest can be formulated as Constraint Satisfaction Problems (CSPs). Solving a general CSP is known to be NP-complete; however, distributed models may take advantage of dividing the problem into a set of simpler interconnected sub-problems which can be more easily solved. In this work, we present a distributed model for solving large-scale CSPs. Our technique carries out a partition over the constraint network by selection of tree structures; after partitioning, the sub-CSPs are arranged into a *meta-tree CSP structure* that is used as a hierarchy of communication by our distributed algorithm. We have focused our research on the railway scheduling problem which can be distributed by tree structures. We show that our distributed algorithm outperforms well-known centralized algorithms.

keywords: Distributed Constraint Satisfaction Problems, Tree Partition, Train Scheduling.

Introduction

Many real problems in Artificial Intelligence (AI) as well as in other areas of computer science and engineering can be efficiently modelled as Constraint Satisfaction Problems (CSPs) and solved using constraint programming techniques. Some examples of such problems include: spatial and temporal planning, qualitative and symbolic reasoning, diagnosis, decision support, scheduling, hardware design and verification, real-time systems and robot planning. Most of the work is focused on general methods for solving CSPs. However, many of the problems solved by using centralized algorithms are inherently distributed. Some works are currently based on distributed CSPs (see special issue of Artificial Intelligence, Volume 161, 2005).

Furthermore, many researchers are working on graph partitioning (Schloegel, Karypis, & Kumar 2003). The main objective of graph partitioning is to divide the graph into a set of regions such that each region has roughly the same

number of nodes and the sum of all edges connecting different regions is minimized. Researchers are almost always interested in the size of nodes or edges, although a few studies have been made on the graph structure of the sub-problems induced by the partition (Miller 1986). For instance, one study seeks a node-separator whose induced graph is hamiltonian). Graph partitioning can also be applied to constraint satisfaction problems. Thus, we can use graph partitioning when dealing with large-scale CSPs to distribute the problem into a set of sub-CSPs. For instance, we can divide a CSP into several subCSPs so that constraints among variables of each subCSP are minimized (Salido & Barber 2006). Otherwise, a domain-dependent partition can be used (see Figure 7). This requires a deeper analysis of the problem to be solved. In this paper, we show that a domain-dependent partition obtains a more adequate distribution so that a higher efficiency is obtained.

In this paper, we present two techniques for structuring and solving binary CSPs. To this end, the binary CSP is previously organized into a *meta-tree CSP structure* so that the original constraint graph is partitioned into trees that represent sub-problems. Thus, the search algorithm carries out the search in each node in linear time (Freuder 1982), (Dechter & Pearl 1987).

Our research is also focused on the railway scheduling problem. Railway traffic has increased considerably, which has created the need to optimize the use of railway infrastructures. This is, however, a hard and difficult task. Our aim is to model the railway scheduling problem as a Constraint Satisfaction Problem (CSP) and solve it using a distributed CSP solver. Due to the topological properties of the railway scheduling problem, the resultant CSP can be distributed in semi-independent sub-problems so that the solution can be solved easier.

In the following section, we summarize some definitions about CSPs. A tree partition method is presented in section 3. Our distributed algorithm for solving *meta-tree CSP structures* is presented in section 4. In section 5 we describe the railway scheduling problem. An evaluation of our methods over real railway networks is presented in section 6. Finally, we summarize our conclusions in section 7.

*This work has been partially supported by the research projects TIN2004-06354-C02-01 (Min. de Educacion y Ciencia, Spain-FEDER), FOM-70022/T05 (Min. de Fomento, Spain), GV/2007/274 (Generalidad Valenciana) and by the Future and Emerging Technologies Unit of EC (IST priority - 6th FP), under contract no. FP6-021235-2 (project ARRIVAL).
Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Centralized, Distributed and Partitionable CSPs

In this section, we present some basic definitions related to CSPs, which will be convenient for our purposes and will unify works from the constraint satisfaction community. Then, we present three ways for solving a CSP: as a centralized problem, as a partitionable problem and as a distributed problem.

A **CSP** consists of: a set of variables $X = \{x_1, x_2, \dots, x_n\}$; each variable $x_i \in X$ has a set D_i of possible values (its domain); a finite collection of constraints $C = \{c_1, c_2, \dots, c_p\}$ that restricts the values that the variables can simultaneously take.

A **solution** to a CSP is an assignment of values to all the variables so that all constraints are satisfied; a problem with a solution is termed *satisfiable* or *consistent*.

A **binary constraint network** is a network in which every constraint subset involves at most two variables. In this case, the network can be associated with a constraint graph, where each node represents a variable and the arcs connect nodes whose variables are explicitly constrained (Dechter 1992).

A **meta-tree CSP structure** is a tree whose nodes are composed by trees. Thus, we will refer to the main tree as *meta-tree CSP structure* and to each individual tree as *single-tree*. We will define each node of the *meta-tree CSP structure* as *meta-node* and each individual and atomic node of the trees as *single-node*. It can be deduced that each *meta-node* corresponds to a *single-tree*. Each constraint between two *single-nodes* of different *meta-nodes* is called **inter-constraint**. Each constraint between two *single-nodes* of the same *meta-node* is called **intra-constraint**.

Partition : A partition of a set C is a set of disjoint subsets of C whose union is C . The subsets are called the blocks of the partition.

Distributed CSP: A distributed CSP (DCSP) is a CSP in which the variables and constraints are distributed among automated agents (Yokoo & Hirayama 2000).

Each agent has some variables and attempts to determine their values. However, there are *inter-constraints* and the value assignment must also satisfy them. In our model, there are k agents $1, 2, \dots, k$. Each agent knows the set of constraints and the domains of variables involved in these constraints.

Partition of CSPs

There are many ways to solve a CSP. However, these problems can be classified into three categories: centralized problems, distributed problems, and partitionable problems.

- A CSP is a *centralized CSP* when there are no privacy/security rules between parts of the problem, and all knowledge about the problem can be gathered into one process. It is commonly recognized that centralized CSPs must be solved by centralized CSP solvers. Many problems are represented as typical examples to be modelled as a centralized CSP and solved using constraint programming techniques. Some examples are: sudoku, n-queens, map coloring, etc.

- A CSP is a *distributed CSP* when the variables, domains and constraints of the underlying network are distributed among agents. This distribution is carried out due to many factors: constraints may be strategic information that should not be revealed to competitors, or even to a central authority; a failure of one agent can be less critical and other agents might be able to find a solution without the failed agent. Examples of such systems are sensor networks, meeting scheduling, web-based applications, etc.
- A CSP is a *partitionable CSP* when the global problem can be divided into smaller problems (sub-problems) which must be coordinated to find the solution to the global problem. For example, the search space of a CSP can be divided into several regions, and a solution is found by using parallel computing.

Given these three categories, we can conclude that a distributed CSP can not be solved by centralized techniques. However, can a centralized CSP be solved by distributed techniques? The answer is 'yes' if the CSP is divided previously.

Real problems usually imply models with a great number of variables and constraints, causing dense networks of inter-relations. Problems of this kind can be handled as a whole only at overwhelming computational cost. Thus, it could be an advantage to decompose problems of this kind into several simpler interconnected sub-problems which can be more easily solved.

In the following example, we show that a centralized CSP could be decomposed into several sub-problems in order to obtain simpler sub-CSPs. In this way, we can apply a distributed technique to solve the decomposed CSP.

The map coloring problem is a typically centralized problem. The goal of a map coloring problem is to color a map so that regions sharing a common border have different colors. Let's suppose that each country of Europe must be colored. Figure 1 (1) shows a colored portion of Europe. This problem can be solved by a centralized CSP solver. However, if the problem is to color each region of each country of Europe (Spain, Figure 1(3); France, Figure 1(4)), it is easy to see that the problem can be partitioned into a set of sub-problems, grouped by clusters. This problem can be solved as a distributed problem, even when the problem is not inherently distributed.

A map coloring problem can be solved by first converting the map into a graph where each region is a vertex and an edge connects two vertices if and only if the corresponding regions share a border. In our problem of coloring the regions of each country of Europe, it can be observed that the corresponding graph maintains clusters representing each country (Spain, Figure 1(3); France, Figure 1(4)). Thus, the problem can be solved in a distributed way.

Why Tree Partition?

As Rina Dechter states in (Dechter 1992), a problem is considered easy when it admits a solution in polynomial time. In the context of constraint networks, a problem is easy if an algorithm like backtracking can solve it in a backtrack-free manner, i.e., without dead-ends, thus producing a so-

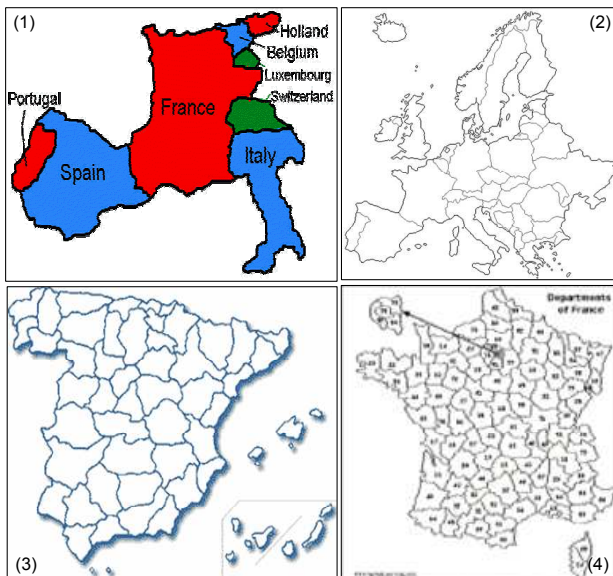


Figure 1: map coloring of Europe.

lution in linear time with regard to the number of variables and constraints. Theoretical researches have identified topological features that determine this level of consistency and have yielded tractable algorithms for transforming some networks into backtrack-free representations. The following paragraphs present a summary of this theory.

The theory is centered on a graphical parameter called *width*, and the definitions are relative to the primal constraint graph. An ordered (primal) constraint graph is defined as one in which the nodes are linearly ordered to reflect the sequence of variable assignments executed by the backtracking algorithm. The *width of a node* is the number of arcs that connect that node to previous ones, the *width of an ordering* is the maximum width of all nodes, and the *width of a graph* is the minimum width of all orderings of that graph. It is known that only trees are width-one graphs (Freuder 1982). An ordered constraint graph is backtrack-free if the level of directional strong consistency along this order is greater than the width of the ordered graph. Thus, if the graph has width-one (i.e., it is a tree), a directional two-consistency is sufficient (Dechter 1992) to solve the problem in linear time.

How to Convert a binary CSP into a Meta-Tree CSP Structure

Any binary CSP can be translated into a *meta-tree CSP structure*. However, there are many ways to divide a graph into trees. Depending on user requirements, it may be desirable to obtain balanced *single-trees*, that is, each *single-tree* maintains roughly the same number of *single-nodes*; or it may be desirable to obtain *single-trees* in such a way that the number of edges connecting two *single-trees* are minimized.

Due to the complexity of finding the best partition, our technique finds an unbalanced tree partition in polynomial time (n^2 in the worst case, where n is the number of vari-

ables). It divides the problem into an undefined number of trees.

The *TreePartition* algorithm (Algorithm 1) divides the network graph G into k sub-graphs which are trees. The nodes and edges of graph G are, respectively, the variables and constraints of the CSP. *TreePartition* randomly selects a root node that does not belong to another sub-graph, then the *SearchTree* function constructs a tree. This function recursively carries out a Depth First Search in graph G . The *SearchTree* function selects a new node i which is connected with *VarNode* and whose inclusion in the present *Tree* does not introduce a cycle, that is, node i is not connected with another node of the present *Tree*. Node i is marked as *visited*; it indicates that this node already belongs to one tree. Tree construction finishes when either there are no *unvisited* nodes or the remainder nodes introduce cycles in the present tree. The *TreePartition* algorithm finishes when all nodes of G belong to any sub-graph.

Algorithm *TreePartition*(G)

Input: Graph G , originally all nodes are unvisited.
Start meta-node v of G

Output: *Tree_partition*

```

Tree_partition =  $\emptyset$ ;
while  $G \neq \emptyset$  do
  Tree =  $\emptyset$ ;
  RootNode = selectNode( $G$ );
  insert RootNode into Tree;
  mark RootNode as visited;
  SearchTree( $G$ , RootNode, Tree);
  insert Tree into Tree_partition;

```

end

end

function SearchTree(G , *VarNode*, *Tree*)

```

forall  $i$  adjacent(1) to VarNode  $\wedge$   $i$  unvisited do
  if NoCycle( $i$ , Tree) then
    insert  $i$  into Tree;
    mark  $i$  as visited;
    SearchTree( $G$ ,  $i$ , Tree);
  end

```

end

```

/* (1) two single-nodes are
   adjacent if at least one
   constraint exists between them.
*/

```

Algorithm 1: *TreePartition* Algorithm.

The next step is to build the *meta-tree CSP structure* with k *meta-nodes* that will be studied by *agents*. This *meta-tree CSP structure* is used as a hierarchy to communicate messages between *meta-nodes*. The *meta-tree CSP structure* is built using Algorithm 2. The nodes and edges of graph G are, respectively, the *meta-nodes* and *inter-constraints* obtained after the CSP partition. The root *meta-node* is obtained by selecting the most constrained *meta-node*. The *MetaTreeCSPStructure* algorithm then simply puts *meta-node* v into the *meta-tree CSP structure* (*process*(v)); it initializes a set of markers to indicate which vertices have

been visited; it chooses a new *meta-node* i and then calls $MetaTreeCSPStructure(G,i)$ recursively. If a *meta-node* has several adjacent *meta-nodes*, it would be equally correct to choose them in any order, but it is very important to delay the test for whether a *meta-node* is visited until the recursive calls for previous *meta-nodes* are finished.

Algorithm $MetaTreeCSPStructure(G,v)$

Input: Graph G , originally all nodes are unvisited.

Start *meta-node* v of G

Output: *meta-tree CSP structure*

$process(v)$;

mark v as visited;

forall *meta-node* i adjacent⁽¹⁾ to v unvisited **do**

$MetaTreeCSPStructure(G,i)$;

end

/ (1) meta-node i is adjacent to meta-node v if at least one inter-constraint exists between i and v . */*

Algorithm 2: $MetaTreeCSPStructure$ Algorithm.

Our aim is to solve CSPs by dividing the constraint graph by means of trees. Figure 3-1 shows a simple example of CSP. In Figure 3-2, this CSP has been divided into several trees and has been translated into a *meta-tree CSP structure*.

DFSTreeSearch Algorithm (DTS)

Our algorithm, called DFSTreeSearch (DTS), can be considered as a distributed and asynchronous algorithm. In the specialized literature, there are many works about distributed CSPs. In (Yokoo & Hirayama 2000), Yokoo et al. present a formalization and algorithms for solving distributed CSPs. These algorithms can be classified as centralized methods, synchronous backtracking or asynchronous backtracking (Yokoo & Hirayama 2000).

DTS is committed to solving the *meta-tree CSP structure* in a *Depth-First Search Tree (DFS Tree)* where the root *meta-node* is composed of the most constrained *single-tree* (in the sense that this *single-tree* maintains a higher number of *single-nodes*). *DFS trees* have already been investigated as a means to boost search (Decher 2003). Due to the relative independence of nodes lying in different branches of the *DFS tree*, it is possible to perform search in parallel on these independent branches.

Once the variables are divided and arranged, the problem can be considered as a distributed CSP, where a group of *agents* manages each *single-tree* with its variables (*single-nodes*) and its constraints (*intra-constraints*). Each *agent* is in charge of solving its own *single-tree* by means of the tree-solving algorithm defined in (Decher 2003). Each sub-problem is composed of its CSP subject to the variable assignment generated by the ancestor *agents* in the *meta-tree CSP structure*.

Thus, the root *agent* works on its sub-problem (root *meta-node*). If the root *agent* finds a solution, then it sends the consistent partial state to its children *agents* in the *meta-tree CSP structure*, and all children work concurrently to solve

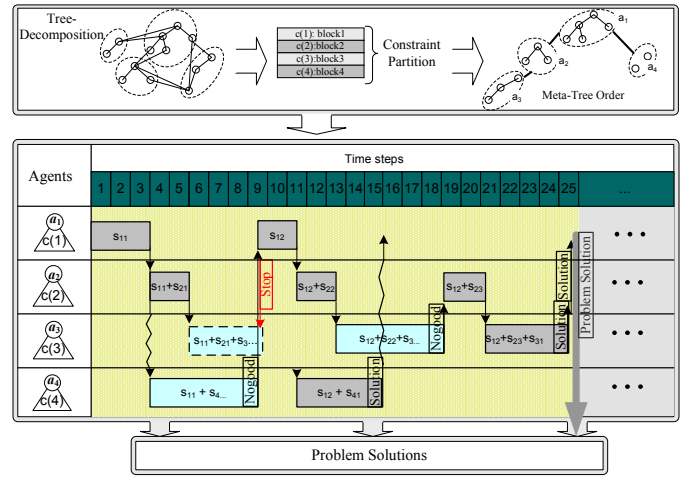


Figure 2: Execution of DFSTreeSearch Algorithm.

their specific sub-problems knowing the consistent partial states assigned by the root *agent*. When a child *agent* finds a consistent partial state, it again sends this partial state to its children and so on. Finally, leaf *agents* try to find a solution to their own sub-problems. If each leaf *agent* finds a consistent partial state, it sends an OK message to its parent *agent*. When all leaf *agents* answer with OK messages to their parents, a solution to the entire problem is found. When a child *agent* does not find a solution, it sends a Nogood message to the parent *agent*. The Nogood message contains the variables that empty the variable domains of the child *agent*. When the parent *agent* receives a Nogood message, it stops the search of the children and tries to find a new solution taking into account the Nogood information and so on. If a parent *agent* finds a new solution, it will start the same process again sending this new solution to its children. Each *agent* works in the same way with its children in the *meta-tree CSP structure*. However, if the root *agent* does not find a solution, then DTS returns *no solution found*.

Figure 2 shows our technique for partitioning a CSP into a *meta-tree CSP structure*. Then, DTS is carried out. Once the CSP is partitioned, the root *agent* (a_1) starts the search process finding a partial solution. It sends this partial solution to its children. *Agents* that are brothers are committed to concurrently finding the partial solutions of their sub-problems. Each *agent* sends the partial problem solutions to its children *agents*. A problem solution is found when all leaf *agents* find their partial solution. For example, (state $s_{12} + s_{41}$) + (state $s_{12} + s_{23} + s_{31}$) is a problem solution. The concurrence can be seen in Figure 2 in *Time step* 4 in which *agents* a_2 and a_4 are concurrently working. *agent* a_4 sends a Nogood message to its parent (*agent* a_1) in step 9 because it does not find a partial solution. Then, *agent* a_1 stops the search process of all its children and it finds a new partial solution which is sent to its children. Now, *agent* a_4 finds its partial solution, and *agent* a_2 works with its child, *agent* a_3 , to find their partial problem solution. When *agent* a_3 finds its partial solution, the global problem will be found. This

happens in *Time step 25*.

Example

Figure 3 shows an example to analyze the behavior of DTS.

First the constraint network of Figure 3(1) is partitioned into three trees and the *DFS tree* is built (Figure 3(2)). Agent *a* finds its first partial solution ($X_1 = 1, X_2 = 1$) and sends it to its children: agents *b* and *c* (see Figure 3(3)). This is a good partial solution for agent *c* (Figure 3(4)); however this partial solution empties the X_3 variable domain. Thus, agent *b* sends a *Nogood* message to its father (*Nogood* ($X_1 = 1$))) (Figure 3(5)). Then, agent *a* processes the *Nogood* message, prunes its search space, finds a new partial solution ($X_1 = 2, X_2 = 2$) and sends it to its children (Figure 3(6)). At this point in the process, agent *c* sends a *Nogood* message to its father (*Nogood* ($X_1 = 2$))) because X_5 variable domain is empty (Figure 3(7)). Agent *a* stops the search of agent *b* (Figure 3(8)) and then processes the *Nogood* message, prunes its search space, finds a new partial solution ($X_1 = 3, X_2 = 3$) and sends it to its children (Figure 3(9)). Since this last partial solution is good for both children, they respond with an *OK* message and the search ends (Figure 3(10)), returning the solution presented in Figure 3(11).




Figure 3: Example of DTS execution.

Railway Scheduling Problem

Train timetabling is a difficult and time-consuming task, particularly in the case of real networks where the number of constraints and the complexity of constraints grow drastically. A feasible train timetable should specify the departure and arrival time of each train to each location of its journey, in such a way that the line capacity and other operational constraints are taken into account. Traditionally, train timetables are generated manually by drawing trains on the

time-distance graph called a running-map. The train schedule is generated from a given starting time and is manually adjusted so that all constraints are met. High priority trains are usually placed first followed by lower priority trains. It can take many days to develop train timetables for a line, and the process usually stops once a feasible timetable has been found. The resulting plan of this procedure may be far from optimal.

A sample of a running map is shown in Figure 4, where several train crossings can be observed. A running map contains information regarding railway topology (stations, tracks, distances between stations, traffic control features, etc.) and the schedules of the trains that use this topology (arrival and departure times of trains at each station, frequency, stops, crossings, etc.). The names of the stations are presented on the left side of Figure 4, and the vertical line represents the number of tracks between stations (one-way or two-way). The horizontal line represents the time.

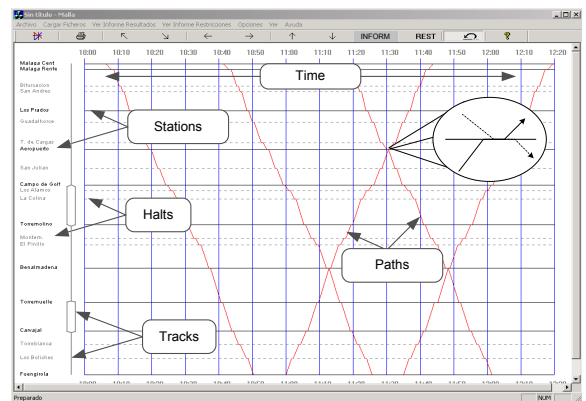


Figure 4: Example of a running-map

The literature of the 1960s, 1970s, and 1980s related to rail optimization was relatively limited. Compared to the airline and bus industries, optimization was generally overlooked in favor of simulation or heuristic-based methods. However, (Cordeau, Toth, & Vigo 1998) point out greater competition, privatization, deregulation, and increasing computer speed as reasons for the more prevalent use of optimization techniques in the railway industry. Our review of the methods and models that have been published indicates that the majority of authors use models that are based on the Periodic Event Scheduling Problem (PESP) introduced by (Serafini & Ukovich 1989). The PESP considers the problem of scheduling as a set of periodically recurring events under periodic time-window constraints. The model generates disjunctive constraints that may cause the exponential growth of the computational complexity of the problem depending on its size. (Schrijver & Steenbeek 1994) have developed CADANS, a constraint programming-based algorithm to find a feasible timetable for a set of PESP constraints. The scenario considered by this tool is different from the scenario that we have been used in this work; therefore, the results are not easily comparable. The train scheduling problem can also be modeled as a special case of

the job-shop scheduling problem ((Silva de Oliveira 2001), (Walker & Ryan 2005)), where train trips are considered jobs that are scheduled on tracks that are regarded as resources.

Our goal is to model the railway scheduling problem as a Constraint Satisfaction Problem (CSP) and to solve it using constraint programming techniques. However, due to the huge number of variables and constraints that this problem generates, a distributed model is developed to distribute the resultant CSP into semi-independent sub-problems so that the solution can be found efficiently.

Variables and Constraints in the Railway Scheduling Problem

The variables of the railway scheduling problem are the arrival and departure times of trains at stations. The variables domain is the time with a granularity of minutes. There are three groups of scheduling rules in our railway scheduling problem: traffic rules, user requirements rules and topological rules. A valid running map must satisfy the above rules. These scheduling rules can be modelled using the following constraints, where variable $TA_{i,k}$ represents that train i arrives at station k and the variable $TD_{i,k}$ means that train i departs from station k :

1. **Traffic rules** guarantee crossing and overtaking operations. We assume two trains (i and j) going in opposite directions between stations k and $k + 1$. The main constraints to take into account are:

- **Reception time constraint.** There exists a given time to detour a train back from the main track so that crossing or overtaking can be performed (RecT).

$$(TA_{i,k} + RecT_i < TA_{j,k}) \vee (TA_{j,k} + RecT_j < TA_{i,k})$$

- **Expedition time constraint.** There exists a given time to put a train back on the main track so that crossing or overtaking can be performed (ExpT).

$$(TD_{i,k} + ExpT_i < TD_{j,k}) \vee (TD_{j,k} + ExpT_j < TD_{i,k})$$

- **Crossing constraint:** Any two trains going in opposite directions must not simultaneously use the same one-way track.

$$(TD_{i,k} + T_{i,(k,k+1)} < TD_{j,k+1}) \wedge (TD_{i,k} < TD_{j,k+1} + T_{j,(k+1,k)})$$

∨

$$(TD_{i,k} + T_{i,(k,k+1)} > TD_{j,k+1}) \wedge (TD_{i,k} > (TD_{j,k+1} + T_{j,(k+1,k)}))$$

- **Overtaking constraint:** Two trains (i and s) going at different speeds in the same direction can only overtake each other at stations.

$$(TD_{i,k} < TD_{s,k}) \wedge (TD_{i,k} + T_{i,(k,k+1)} < TD_{s,k} + T_{s,(k,k+1)})$$

∨

$$(TD_{i,k} > TD_{j,k}) \wedge (TD_{i,k} + T_{i,(k,k+1)} > (TD_{s,k} + T_{s,(k,k+1)}))$$

2. **User Requirements:** The main constraints due to user requirements are:

- **Type and Number of trains** going in each direction to be scheduled.
- **Path of trains:** Locations used and **Stop time** for commercial purposes in each direction.

$$TD_{i,k} = TA_{i,k} + StopTime_{i,k}$$




Figure 5: Constraints related to crossing and overtaking in stations

- **Scheduling frequency.** Train departure must satisfy frequency requirements in both directions. It could be a fixed time (1) or a time interval ($Freq \pm \delta$) (2). Frequency is a very tight constraint and is only sometimes required.

$$(1) TD_{i+1,k} = TD_{i,k} + Freq$$

$$(2) (TD_{i,k} + Freq - \delta) \leq TD_{i+1,k} \leq (TD_{i,k} + Freq + \delta)$$

- **Departure interval** for the departure of the first trains going in both the up and down directions.

$$StartTime_i < TD_{i,1} < EndTime_i$$

3. **Railway infrastructure topology and type of trains** to be scheduled give rise to other constraints to be taken into account. Some of them are:

- **Number of tracks in stations** (to perform technical and/or commercial operations) and the **number of tracks between two locations** (one-way or two-way).
- **Time constraints**, between each two contiguous stations ($T_{i,(k,k+1)}$).

$$TA_{i,k+1} - TD_{i,k} = T_{i,(k,k+1)}$$

Figure 6 shows the set of variables of two trains going in opposite directions between stations A and E . After studying the problem, we have detected an advantageous tree partition of the railway scheduling problem. In Figure 6, the edges among two variables represent time constraints ($TD_{i,k} - TA_{i,k+1}$) and stop time constraints ($TA_{i,k} - TD_{i,k}$) respectively; these constraints make up train paths and they could be private information of railway operators. Figure 7 shows a clear tree partition of the railway scheduling problem where each sub-CSP has all variables of only one train and the respective *intra-constraints* are time constraints and time stop constraints that are usually fixed by railway operators. The *inter-constraints* are the traffic rules that are usually controlled by infrastructure managers.

Evaluation

In this section, we carry out an evaluation between DTS and a centralized CSP solver. To this end, we have used a well-known centralized CSP solver called *Forward Checking* (FC). The classical binary version of FC ((Haralick & G. 1980)) has a prohibitive computational cost for the type of problems evaluated in this section (a simple railway scheduling problem with 4 trains could not be solved after 1 day of execution); that is why we use the *full path consistency Forward Checking* algorithm (FCPath)¹. This algorithm per-

¹FCPath were obtained from Van Beek page. It can be found in: <http://ai.uwaterloo.ca/~vanbeek/software/software.html>

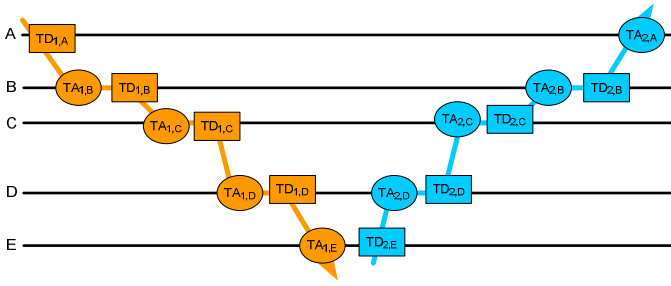


Figure 6: Variables of two opposite trains.

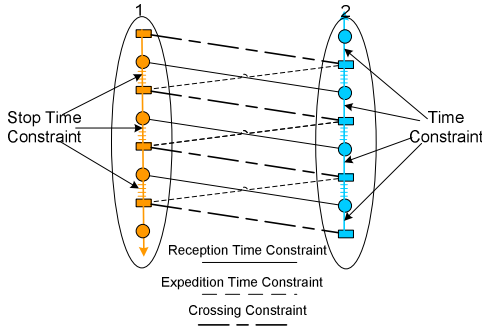


Figure 7: Train-based tree partition.

forms full path consistency on future variables whenever the current variable is instantiated.

This empirical evaluation of the railway scheduling problem was carried out over a real railway infrastructure that joins two important Spanish cities (La Coruna and Vigo). The journey between these two cities is currently divided by 40 stations. In our empirical evaluation, each set of random instances was defined by the 3-tuple $\langle n; s; f \rangle$, where n was the number of periodic trains in each direction, s the number of stations, and f the frequency (in minutes). The problems were randomly generated by modifying these parameters. Usually, increasing the number of trains involves a CSP with a greater number of variables; increasing the number of stations involves a CSP with a greater number of variables and a greater domain size; and decreasing the frequency involves increasing the problem tightness because the number of conflicts between trains is greater.

In this evaluation, we compare the running time and *concurrent constraint checks* (CCC) ((Meisels *et al.* 2002)) of DTS with a well-known centralized algorithm: FCPATH. DTS is executed over two different tree partitions: *random partition* is obtained using the *TreePartition* algorithm (Algorithm 1), which is a general tree partition method; *train partition* is a domain-dependent partition, which has been illustrated in Figure 7. It must be taken into account that both types of partitions roughly generate the same number of sub-problems and *inter-constraints*.

Figures 8 and 9 show the behaviors of DTS and FCPATH in several instances of n according to the tuple $\langle n, 5, 60 \rangle$. The number of trains (n) was increased from 1 to 20 trains in each direction. It must be taken into account that both




Figure 8: Running Times in problems $\langle n, 5, 60 \rangle$.

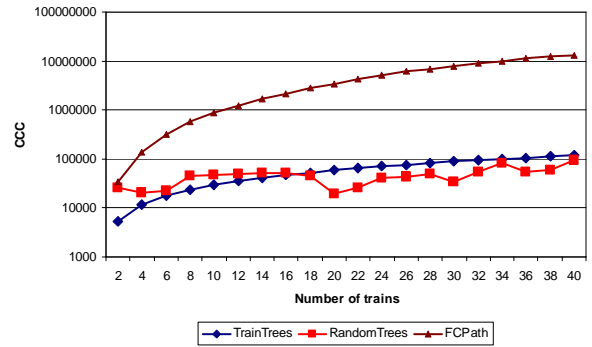


Figure 9: Concurrent constraint checks in problems $\langle n, 5, 60 \rangle$.

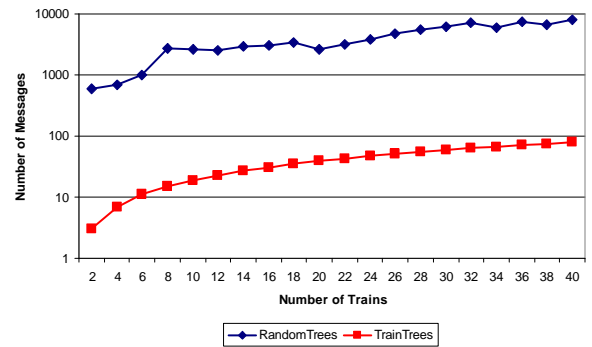


Figure 10: Message exchanges in problems $\langle n, 5, 60 \rangle$.

graphs maintain a \log_{10} scale. Figures 8 and 9 show that DTS outperforms the FCPATH algorithm, both *random partition* and *train partition*, in all instances. Figure 8 shows that DTS with *train partition* always has smaller running times than DTS with *random partition*. However, Figure 9 shows that DTS with *random partition* sometimes has fewer CCC than DTS with *train partition*. The two last assertions seem to be contradictory, but the explanation is in Figure 10: DTS with *train partition* exchanges fewer messages than DTS with *random partition*; thus, *train partition* saves a lot of running time. This is due to *train partition* involves better agent coordination.

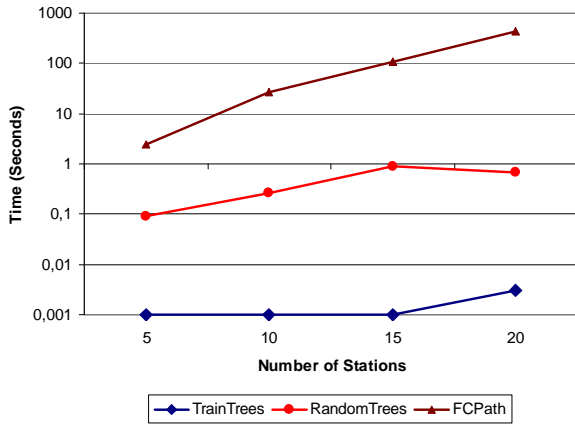


Figure 11: Running Times in problems $\langle 4, s, 60 \rangle$.




Figure 12: Running Times in problems $\langle 4, 10, f \rangle$.

Figure 11 show the behaviors of DTS and FCPath in several instances of s according to the tuple $\langle 4, s, 60 \rangle$, where the number of stations (s) was increased from 5 to 20. It can be observed that DTS with *train partition* has a better behavior than DTS with *random partition* and FCPath. In Figure 12, we evaluate the behavior of the algorithms with different frequencies according to the tuple $\langle 4, 10, f \rangle$, where frequency was increased from 15 to 60 minutes. It can be observed that due to the problem instances maintain the same number of variables and domain size, DTS with *train partition* and FCPath maintain homogeneous behaviors. In general, both graphs corroborate the good behavior of the DTS algorithm, particularly with *train partition*.

Conclusions

We have presented two techniques for structuring and solving binary CSPs. The first one translates, in polynomial time, the original binary graph into a *meta-tree CSP structure*, where each node in the *meta-tree CSP structure* is a tree. The second technique is a distributed algorithm (DTS) for solving the resultant *meta-tree CSP structure*. DTS exploits the linear complexity to solve each tree and minimizes the storage of Nogoods. These techniques have been ap-

plied to the railway scheduling problem. The evaluation shows that general distributed models have a better behavior than the centralized model, but domain-dependent distributed models are more efficient than general ones. Thus, this technique may be appropriate for solving centralized problems that can be divided in smaller sub-problems.

References

- Cordeau, J.; Toth, P.; and Vigo, D. 1998. A survey of optimization models for train routing and scheduling. *Transportation Science* 32:380–446.
- Decher, R. 2003. *Constraint Processing*. Morgan Kaufmann.
- Dechter, R., and Pearl, J. 1987. Network-based heuristics for constraint satisfaction problems. *Artificial Intelligence* 34:1–38.
- Dechter, R. 1992. Constraint networks (survey). *Encyclopedia Artificial Intelligence* 276–285.
- Freuder, E. 1982. A sufficient condition for backtrack-free search. *Journal of the ACM* 29:24–32.
- Haralick, R., and G., E. 1980. Increasing tree efficiency for constraint satisfaction problems. *Artificial Intelligence* 14:263–314.
- Meisels, A.; Kaplansky, E.; Razgon, I.; and Zivan, R. 2002. Comparing performance of distributed constraint processing algorithms. In *Proc. 4th Workshop on Distributed Constraint Reasoning*.
- Miller, G. 1986. Finding small simple cycle separators for 2-connected planar graphs. *Journal of Computer and System Sciences* 32:265–279.
- Salido, M., and Barber, F. 2006. Distributed csps by graph partitioning. *Applied Mathematics and Computation* 183:491–498.
- Schloegel, K.; Karypis, G.; and Kumar, V. 2003. Graph partitioning for high-performance scientific simulations. *Sourcebook of parallel computing* 491–541.
- Schrijver, A., and Steenbeek, A. 1994. Timetable construction for railned. *Technical Report, CWI, Amsterdam, The Netherlands*.
- Serafini, P., and Ukovich, W. 1989. A mathematical model for periodic scheduling problems. *SIAM Journal on Discrete Mathematics* 550–581.
- Silva de Oliveira, E. 2001. Solving single-track railway scheduling problem using constraint programming. *Phd Thesis. Univ. of Leeds, School of Computing*.
- Walker, C., S. J., and Ryan, D. 2005. Simultaneous disruption recovery of a train timetable and crew roster in real time. *Comput. Oper. Res* 2077–2094.
- Yokoo, M., and Hirayama, K. 2000. Algorithms for distributed constraint satisfaction: A review. *Autonomous Agents and Multi-Agent Systems* 3:185–207.