

Feasible Distributed CSP Models for Scheduling Problems*

Miguel A. Salido

Department of Information Systems and Computation
Technical University of Valencia, Spain
msalido@dsic.upv.es

Abstract

Nowadays, many real problems can be formalized as Distributed CSPs. A distributed constraint satisfaction problem (DisCSP) is a CSP in which variables and constraints are distributed among multiple automated agents. Many researchers assume for simplicity that each agent has exactly one variable. For real planning and scheduling problems, these distributed techniques require a large amount of messages passed among agents, so these problems are very difficult to solve. In this paper, we present a general distributed model for solving real-life scheduling problems and propose some guidelines for distributing large-scale problems. Furthermore, we present two case studies in which two scheduling problems are distributed by using our model.

Introduction

In recent years we have seen an increasing interest in Distributed Constraint Satisfaction Problem (DisCSP) formulations to model combinatorial problems (see the special issue on Distributed Constraint Satisfaction in the Artificial Intelligence journal, vol 161, 2005). There is a rich set of real-world distributed applications, such as network systems, planning, scheduling, resource allocation, etc, for which the DisCSP paradigm is particularly useful. In such distributed applications, privacy issues, knowledge transfer costs, robustness against failure, etc preclude the adoption of a centralized approach (Faltings & Yokoo 2005).

Briefly, 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 $C = \{c_1, c_2, \dots, c_p\}$ 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*.

*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, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

A distributed CSP is a CSP in which the variables and constraints are distributed among automated agents (Yokoo & Hirayama 2000). Finding a value assignment to variables that satisfies inter-agent constraints can be viewed as achieving coherence or consistency among agents.

The most cited papers related to DisCSP make the following assumptions for simplicity in describing the algorithms:

1. Each agent has exactly one variable.
2. All constraints are binary.
3. Each agent knows all constraint predicates relevant to its variable.

Although the great majority of real problems are naturally modelled as non-binary CSPs, the second assumption is comprehensible due to the fact that there exist some techniques that translate any non-binary CSP into an equivalent binary one (Bacchus & van Beek 1998).

However, the first assumption is too restrictive and the main basic research focuses on small instances. Also, little work has been done to solve real-life problems.

Main Features in DisCSPs

If all knowledge about the problem could be gathered into one agent, this agent could solve the problem alone using traditional centralized constraint satisfaction algorithms. However, such a centralized solution is often inadequate or even impossible. Faltings and Yokoo (Faltings & Yokoo 2005) present some reasons why distributed methods may be desirable:

- The cost of creating a central authority. A CSP may be naturally distributed among a set of agents. In such cases, a central authority for solving the problem would require adding an additional element that was not present in the architecture. Examples of such systems are sensor networks or meeting scheduling.
- The knowledge transfer costs: In many cases, constraints arise from complex decision processes that are internal to an agent and cannot be articulated to a central authority. Examples of this range from simple meeting scheduling, where each participant has complex preferences that are hard to articulate, to coordination decisions in virtual enterprises that result from complex internal planning. A

centralized solver would require such constraints to be completely articulated for all possible situations. This would entail prohibitive costs.

- Privacy/Security concerns: Agents involve constraints that may represent strategic information that should not be revealed to competitors, or even to a central authority. This situation often arises in many enterprises. Privacy is easier to maintain in distributed solvers.
- Robustness against failure: The failure of the centralized server can be fatal. In a distributed method, a failure of one agent can be less critical and other agents might be able to find a solution without the failed agent. Such concerns arise, for example, in sensor networks, but also in web-based applications where participants may leave while a constraint solving process is ongoing.

These reasons have motivated significant research activity in distributed constraint satisfaction. Up to now, the field has reached a certain maturity and has developed a range of different techniques. Nevertheless, most of the works are focused on developing new techniques which are evaluated using toy problems and random benchmarks.

Open Issues in DisCSPs

In spite of significant progress, there are many important open issues in distributed CSP. The five main open issues for using distributed CSP are the followings:

- While distributed algorithms eliminate the need for a central authority, the currently known algorithms pay a high price in efficiency. In general, the message traffic even for a single agent can be higher than what would be required to communicate the entire problem to a leader agent that could solve it centrally. More research is required to significantly reduce the communication requirements, possibly with radically different algorithms that are better suited for distribution.
- Many DisCSP algorithms assumes an agent has enough knowledge to evaluate constraints that are related to its variables. If this is not true, some constraints may still have to be communicated or additional communication may be needed. Also, more research is needed on algorithms that minimize the number of constraint evaluations when evaluating constraints is costly.
- While there are algorithms using cryptographic techniques to ensure complete privacy of agent constraints, their message complexity is very high. For most other DisCSP algorithms, there is no good characterization of how much information is revealed to other agents. More research is needed on measures of privacy loss and on algorithms that balance the trade-off between privacy loss and efficiency.
- While most distributed algorithms tolerate certain kinds of agent failures, there is no good characterization of the kind of failures that are allowed for each algorithm. In general, this issue has not yet been given significant attention in research.

- While most distributed algorithms manage only one variable per agent, there are no specific distributed algorithms to solve real-world problems in an efficient way. This issue has not yet been given enough attention in research. This paper focuses on this issue.

Other issues that are important but have received little attention so far include openness, i.e., the possibility to add and remove agents dynamically during execution, and incentive-compatibility, i.e., making algorithms safe against manipulation by self interested agents.

From Basic Research Toward Applied Research

One of the pioneer researchers in DisCSP said "*So far, we assume that each agent has only one local variable. Although the developed algorithms can be applied to the situation where one agent has multiple local variables by the following methods, both methods are neither efficient nor scalable to large problems*" (Yokoo & Hirayama 2000).

- Method 1: each agent finds all solutions to its local problem first. By finding all solutions, the given problem can be re-formalized as a distributed CSP, in which each agent has one local variable whose domain is a set of obtained local solutions. Then, agents can apply algorithms for the case of a single local variable. The drawback of this method is that when a local problem becomes large and complex, finding all the solutions of a local problem becomes virtually impossible.
- Method 2: an agent creates multiple virtual agents, each of which corresponds to one local variable, and simulates the activities of these virtual agents. However, since communicating with other agents is usually more expensive than performing local computations, it is wasteful to simulate the activities of multiple virtual agents without distinguishing the communications between virtual agents within a single real agent, and the communications between real agents.

In spite of significant progress in distributed CSP, the following question is straightforward: *Why is only one variable per agent assumed?* (Salido 2007). Only some works include a set of variables into an agent (Silaghi & Faltings 2005), (Ezzahir & Bouyakhf 2007), (Salido & Barber 2006). Nevertheless, few works have been focused on distributed techniques for solving large scale problems (Yokoo *et al.* 1998). In this paper, we present different alternatives for managing large scale problems. Each agent will be committed to a large number of variables and constraints and several subproblems can be executed concurrently depending on the internal structure.

A General Distributed Model

Depending on the problem to be considered, the distributed model will maintain different properties. Our general model for solving scheduling problems can be considered as a synchronous model. It is meant to be a framework for interacting holons/agents to achieve a consistent state. The main idea of our holonic/multi-agent system model is based on

(Salido, Giret, & Barber 2003) in which the problem is partitioned into a set of subproblems; then the subproblems are classified in the appropriate order and are solved concurrently. In this section, we introduce the notion of holon as a complementary idea of agent. Depending on the scheduling problem, we will use holons instead of agents. A holon is an autonomous and cooperative unit that can be seen as a whole and a part (Koestler 1971). Therefore, a holarchy is a group of basic holons and/or recursive holons that are themselves holarchies. A Holonic architecture (HMS 1994) (Koestler 1971) is committed to organizing entities (holon or agent (Giret & Botti 2004)) that are responsible for solving each subproblem.

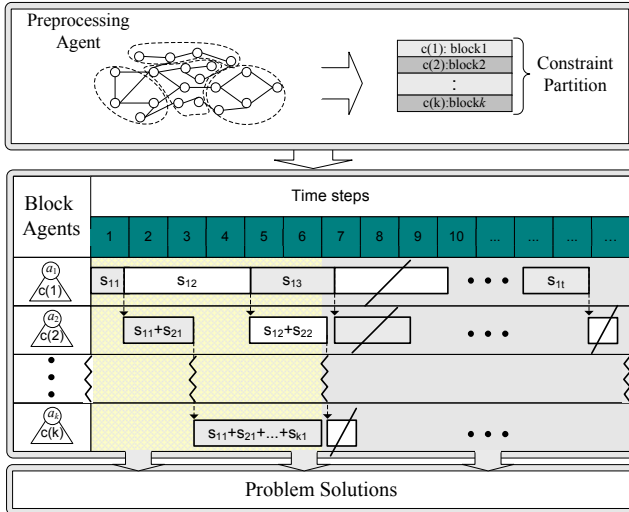


Figure 1: General Distributed Model.

Depending on the problem, it is partitioned in k blocks or clusters in order to be studied by holons/agents called *block agents*. Furthermore, a partition agent is committed to classifying the subproblems in the appropriate order depending on the selected proposal.

Once the problem is divided into k blocks by a *preprocessing agent*, a group of *block agents* concurrently manages each block of constraints. Each *block agent/holon* is in charge of solving its own subproblem by means of a search algorithm. Each *block agent/holon* is free to select any algorithm to find a consistent partial state. It can select a local search algorithm, a backtracking-based algorithm, or any other, depending on the problem topology. In any case, each *block agent/holon* is committed to finding a solution to its particular subproblem. This subproblem is composed by its CSP subject to the variable assignment generated by the previous *block agents/holons*. Thus, *block agent/holon 1* works on its group of constraints. If *block agent/holon 1* finds a solution to its subproblem, then it sends the consistent partial state to *block agent/holon 2*, and they both work concurrently to solve their specific subproblems; *block agent/holon 1* tries to find another solution and *block agent/holon 2* tries to solve its subproblem knowing that its *common variables* have been assigned by *block agent/holon 1*. This second so-

lution found by *block agent/holon 1* is stored to be sent to *block agent/holon 2* if it was necessary. In this case, *block agent/holon 2* will not wait for *block agent/holon 1* to search for a new solution.

Thus, *block agent j*, with the variable assignments generated by the previous *block agents/holons*, works concurrently with the previous *block agents/holons* and tries to find a more complete consistent state using a search algorithm. Finally, the last *block agent/holon k*, working concurrently with *block agents/holons 1, 2, ..., (k - 1)*, tries to find a consistent state in order to find a problem solution.

Figure 1 shows the holonic system, in which the *preprocessing agent* carries out the network partition and the *block agents/holons* (a_i) are committed to concurrently finding partial problem solutions (s_{ij} , where i denotes the number of *block agents/holons* and j the j th solution). Each *block agent/holon* sends the partial problem solutions to the following *block agent/holon* until a problem solution is found (by the last *block agent/holon*). For example, state: $s_{11} + s_{21} + \dots + s_{k1}$ is a problem solution. The concurrence can be seen in Figure 1 in *Time step 6* in which all *block agents/holons* are concurrently working. Each *block agent/holon* maintains the corresponding domains for its *new variables*. The *new variables* are the variables that are not involved in previous *block agent/holon*. The *block agent/holon* must assign values to its *new variables* so that the block of constraints is satisfied. When a *block agent/holon* finds a value for each *new variable*, it then sends the consistent partial state to the next *block agent/holon*. When the last *block agent/holon* assigns values to its *new variables* satisfying its block of constraints, then a solution is found. It must be taken into account that if a *block agent/holon* maintains too many variables or constraints, it can be decomposed into a set of new agents/holons.

Some Guidelines for Distributing Large-Scale Problems

To solve large-scale problems, we must distribute the problems taking into account some guidelines:

1. *The number of subproblems (agents)*. They must be in concordance with the size of the problem. As we have pointed out in the previous sections, one agent per variable is unmanageable. A problem with thousands of variables and constraints cannot be modelled as a distributed model with thousands of agents due to the high computational cost in the solving process. Generally and due to privacy issues, the number of subproblems is straightforward, given by the nature of the problems. Nevertheless, some subproblems are too large and they can be divided/decomposed again into smaller ones in order to be solved in a reasonable time. A reasonable way to divide the problem is by means of graph partitioning techniques (Salido & Barber 2006). However, in many real problems, the best way to partition the problem is by carrying out a domain dependent partition. Following, we present an example of railway scheduling problem distributed by types of trains and sets of stations.
2. *The order in which each subproblem is executed (agent*

priority). Sometimes all subproblems can be executed concurrently and partial states are sent to their neighbours. In many other cases, the subproblems can be ordered using a selected criteria.

- From the subproblem with the most neighbours to the subproblems with the least neighbours.
- From the tightest subproblems to the loosest ones.
- From the subproblems that maintains hard constraints to the subproblems with soft constraints.
- etc...

This ordering does not mean that the subproblems are solved sequentially, but rather that some subproblems are executed first and then all neighbours are concurrently executed (see Figure 1). Furthermore, this ordering, represented by priorities among agents, can change dynamically depending on many factors such as number of no-goods, number of backtracks, etc.

3. *The management of backtracking.* When an agent does not find a partial solution, it must communicate its current state to the related agents in order to avoid unnecessary searches. Some works decompose the subproblem into a set of subproblems in such a way that the resultant problem is represented as a tree. Each node in the tree is a subproblem (Abril, Salido, & Barber 2007). Once a node/agent finds a partial solution, it is sent to its children, and they are solved concurrently. However, if a node/agent does not find a partial solution, a message is sent to its parents, and, depending on the management of the backtracking (which parent it backtracks), the search space will be pruned more efficiently.
4. *The necessity of a central authority.* As we have pointed out, depending on the problem, a central regulatory authority will be necessary or not. Many researchers suggest, that for some multi-agent systems, no central regulatory authority is needed and can be replaced by a virtual representation where each agent is responsible for maintaining its own partial view of the relevant institutional state. This is due to the fact that, if the central authority fails, the problem cannot be solved. However, many real problems are hierarchical by nature and it is necessary to generate different levels of abstraction. In the next section, we will present a global road transportation system: a hierarchical system that maintains as many levels as necessary, depending on the problem complexity.

Distributed Models for Distributed Scheduling Problems: Two case studies

Many real problems are distributed by nature. However, this distribution does not imply one variable per agent. For instance, many scheduling problems are decentralized by nature, and the problem is decomposed in clusters. Each cluster (composed by variables and constraints) is solved by an agent, and it communicates a consistent partial state. In this way, a large-scale scheduling problem can be solved with reasonable efficiency, maintaining all privacy issues. Most

of these problems are domain-dependent and general distributed models (as the model proposed above) are not appropriate. Therefore, domain dependent distributed models must be developed to efficiently manage these problems.

In this section, we present two real-life scheduling problems, which are very complex problems and are distributed by nature. To solve them in a distributed way, we group several compatible variables per agent so that the problem can be solved in a reasonable time.

Railway Scheduling Problem

Train timetabling is a difficult problem, 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. 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.

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 et al. (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 and Ukovich (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 and Steenbeek (Schrijver & Steenbeek 1994) have developed CADANS, a constraint programming-based algorithm to find a feasible timetable for a set of PESP constraints. The train scheduling problem can also be modeled as a special case of the job-shop scheduling problem (Silva de Oliveira (Silva de Oliveira 2001), Walker et al. (Walker & Ryan 2005)), where train trips are considered as jobs that are scheduled on tracks that are regarded as resources.

We modelled the railway scheduling problem as a Constraint Satisfaction Problem (CSP) and it was solved by using constraint programming techniques. However, due to the distributed nature of the problem and the huge number of variables and constraints that this problem generates, a distributed model was developed to distribute the resultant CSP into semi-independent subproblems so that the solution can be found efficiently (Salido *et al.* 2007).

Here, we present two ways for distributing the railway scheduling problem. It is partitioned into a set of subproblems by means of types of trains and by means of contiguous constraints.




Figure 2: Distributed Railway Scheduling Problem by type of trains.

Distribution by types of trains This distributed model is based on dividing the original railway problem by means of train types. In this model, each agent is committed to assigning values to variables related to a train or sets of trains in order to minimize the journey time. This partition model takes into account some of the guidelines given above.

- Depending on the problem instance, the number of partitions will be given by the railway operator or by the number of trains. Figure 2 shows a running map with 20 partitions. Each agent manages one train. Each train generates a large number of variables and constraints, depending on the number of stations and the user requirements. Furthermore, this model allows us to improve privacy. Currently, due to the policy of deregulation in the European railways, trains from different operators work on the same railway infrastructure. In this way, the partition model also gives us the possibility of partition the problem so that each agent is committed to an operator. Thus, different operators maintain privacy about strategic data.
- This model allows us to efficiently manage priorities between different types of trains (regional trains, high-speed trains, freight trains). In this way, agents committed to priority trains (high-speed trains) will first carry out value assignment to variables in order to achieve better journey times. This ordering is inserted into our distributed model to solve the scheduling problem concurrently.

Distribution by set of stations This distributed model is based on distributing the original railway problem by means of contiguous stations. The deregulation of European railway operators gives the opportunity to schedule long journeys. However, long journeys involve large number of stations in different countries with different railway policies. Therefore, a logical partition of the railway network can be

carried out by means of regions (contiguous stations). Some of the guidelines presented above that must be taken into account in this model are:

- The number of subproblems depends directly on several factors: distance of the journey, number of different regions (mainly countries) and railway topology. This distributed model divides the problem into a set of physical regions. It is important to analyze the railway infrastructure in order to detect restricted regions (bottlenecks). To balance the problem, each agent is committed to a different number of stations. An agent can manage many stations if they are not restricted stations; however an agent can manage only a few stations if they represent bottlenecks.




Figure 3: Distributed Railway Scheduling Problem by set of stations.

- The order in which each subproblem is executed plays an important role. Agents committed with bottlenecks have preference to assign values to variables due to the fact that their domains are reduced (variable ordering). Once the agent with the tightest constraints solves its subproblems, both the previous agent and the following agent must concurrently solve their own subproblems.
- The management of backtracking is very important to avoid unnecessary constraint checking. For instance, if agent 4 in Figure 3 finds a partial solution to its subproblem, it communicates to agent 3 and agent 5. Both agents concurrently search to find their own partial solutions. However, if agent 3 does not find a partial solution, it sends a nogood message to agent 4, and this last agent sends a message to agent 5 in order to stop its search. As we pointed out in Figure 1, while agent 3 and agent 5 work concurrently to find their partial solutions, agent 4 also works in the same time step in order to find another partial solution. This last partial solution will be used if agent 3 or agent 5 backtracks due to inconsistency with the previous partial solution given by agent 4.

Figure 3 shows the journeys to be scheduled between two cities. They are decomposed into several shorter journeys. The set of stations are partitioned in blocks of contiguous

stations and a set of agents will coordinate with each other to achieve a global solution. Thus, we can obtain important results such as railway capacity (Abril *et al.* 2007), consistent timetable, etc.

Global Road Transportation System

During recent years the development of automated traffic systems has received increased attention, and substantial effort has been invested in trying to find a solution to problems associated with road transport. Among these problems are road accidents caused by human-related factors, such as tiredness, loss of control, a slow reaction time, limited field of view, etc. A further transport-related problem is that of loss of time which may be caused by slow driving speed due to weather conditions, road conditions, visibility, and traffic congestion. In this section, we present a global road transportation system, which is being developed by several European Universities. The main goal of the algorithmic section is to develop algorithms capable of creating driving schemes for a vehicle from any arbitrary address to any other arbitrary address (in the address space of the system), while considering, and if necessary adapting, the driving schemes of other vehicles travelling in the system at the same time according to priorities, driving and optimization rules. Thus, distributed techniques are necessary for solving these problems.

The Global Automated Transportation System (GATS) (Zelinkovskyn) is a driverless, integrated transport system. It has the astonishing ability to simultaneously coordinate the macro and micro needs of road transport networks. Millions of vehicles can be optimally, simultaneously and automatically "driven" over a virtually unlimited geographic region, including whole continents, while the requirements of each individual vehicle and its passengers are attended to at the same time. It is an innovative concept, based on simple, recognized principles and proven technologies.

Its application will revolutionize road travel by dramatically increasing safety, reducing congestion, and eliminating driving-associated stress and fatigue. The consequence...an overall improvement in the quality of everyday life.

Due to the decentralized and modular nature of the architecture it can be implemented with the same ease and simplicity in both small contained areas such as airports and theme parks and in larger areas such as local, national and international road systems.




Figure 4: Driving on the System.

Following, we summarize the architecture above and below the road. In the center of a traffic lane, 15-20 cm below the road surface there is an "Intelligent Cable" of about 1 cm in diameter which is comprised of tiny intelligent transponders (Road-Units (RUs)) located at fixed distances (less than a vehicle length of 3m) from each other. While driving, the

vehicle sends short radio transmissions down towards the RUs at regular time intervals (about every 30 milliseconds). The RU receives a transmission, processes it, and responds with a radio transmission back to the vehicle. The vehicle communicates continuously with the RUs one after the other incessantly. Thus, it has continuous radio communication with the RUs and the whole system connected to them. The interchange of radio transmissions between the vehicles and the RUs also facilitates lateral and longitudinal positioning of vehicles on the road, as presented in Figure 4.

The memory of each RU stores the specifications of the RU and individual driving instructions that it will transmit to each vehicle above it. Several hundreds of consecutive RUs constitute a Segment, whose functions are administered by a *Segment Controller*. The *Segment Controller* is connected to its RUs through the Parallel Buses and is responsible for "driving" the vehicles passing in its domain; performing routine maintenance check-ups of the components in its segment; and monitoring and regulating their mutual performance. A group of adjacent *Segment Controllers* has a superior controller, the *Level-1 Controller*, which coordinates and controls its individual and mutual functions. A group of adjacent *Level-1 Controllers* has a superior controller: *Level-2 Controller*. This goes on hierarchically (Figure 5). The *Top Level Controller* coordinates and controls the functions of the whole system. There is virtually no limit to the number of levels and to the size of the geographic domain of the systems.

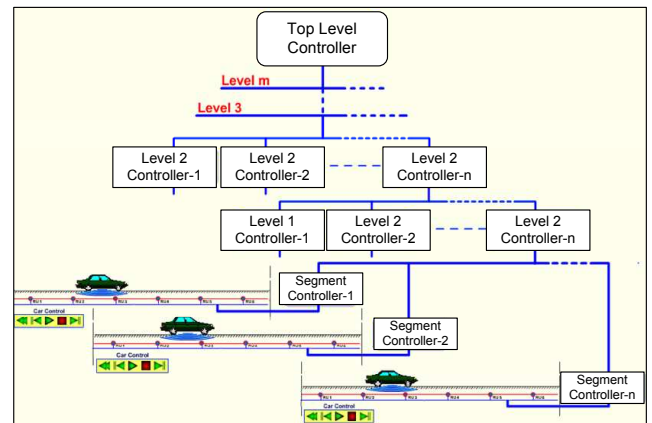


Figure 5: Hierarchical Architecture of the System.

Integrated Functioning Assume a vehicle is in a parking lot above a RU. The passengers turn the vehicle on, which begins to send short radio transmissions down towards the road. The RU detects those transmissions and responds with radio transmissions back to the vehicle. The RU initiates a communication session with the *Segment Controller* in order to inform it about the new event. The passengers in the vehicle enter their requirements as destination, priority, preferred routes etc. The vehicle's processor sends a message to the *Segment Controller* which includes the requirements, the exact location of the vehicle relative to the RU and its own specifications. The *Segment Controller* processes the

request while considering additional inputs from other RUs in the Segment and from its superior Controller. Finally it prepares a driving instruction message for each RU in the Segment. The RUs will send these instructions to the vehicle when it passes above them. Each message includes an addressee (RU1 etc.), a vehicle ID, the expected arrival time of the vehicle to the RU, the speed that the vehicle should travel at and the driving direction. When the vehicle is driving from one RU to another, the active RU uses the Serial Bus to inform the next and previous RUs in the sequence about the exact timing, the ID number and other specifications of the moving vehicle. If the RUs detect intolerable deviation from the plan, they can initiate a so-called Emergency Braking Procedure. The active RU uses the Parallel Bus to inform the *Segment Controller* with information about the moving vehicle.

A Distributed Model for GATS Traditional Centralized techniques fail to model and implement problems of this type due to their complex and large nature. Due to the decentralized and modular nature of the architecture, the algorithms to obtain the scheduling of each vehicle must be distributed. Figure 6 shows the map of Europe to be distributed/divided into regions (countries); each region is divided into sub-regions, and so on.



Figure 6: Map of Europe to be distributed.

Briefly, the system is composed by a network, where nodes are locations and arcs are roads. Depending on the granularity, nodes are points in the road or regions in a country. In the lower level of the system, each RU is represented by a variable (see Figure 7). The system may be composed of millions of RUs. As we have explained in the first section, this problem cannot be managed by current distributed CP techniques by using a variable per agent. By using these approaches, the problem generates millions of agents and messages passing in the interaction scenarios. This makes the resulting DisCSP unmanageable.

To overcome these weakness, we use our distributed

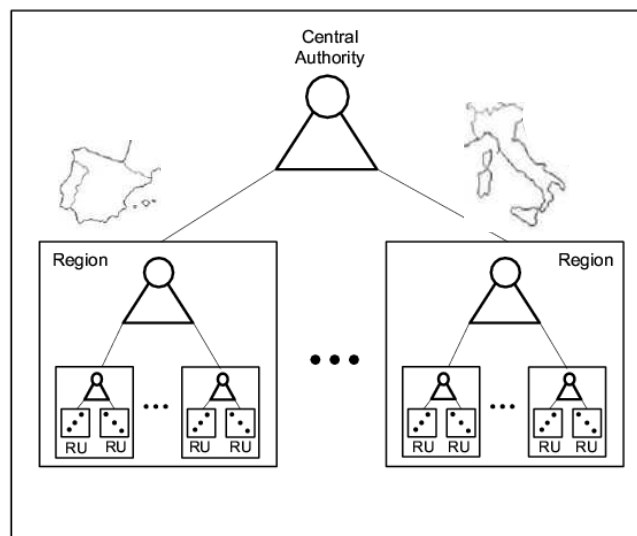


Figure 7: Distributed model with a central authority.

model in which the problem is partitioned into subproblems that represent regions, countries, etc (see Figure 7). Here, we use a Holonic architecture (HMS 1994)(Koestler 1971) to organize the entities (holon or agent (Giret & Botti 2004)) that are responsible for solving each subproblem. As we have pointed out above, a holon is an autonomous and cooperative unit that can be seen as a whole and a part (Koestler 1971).

The distributed model generated for this scheduling problem follows the guidelines presented in the paper.

- The number of subproblems depends on the size of the system. A holon can represent a track between two traffic lights or represent a region or a country. Figure 7 shows two holons that represent two countries, Spain and Italy. Each of them is composed of a set of sub-holons that represent regions, and each sub-holon is composed by new sub-holons that represent sub-regions and so on. The base case is composed of individual variables that represent RUs.
- The execution of the subproblems is carried out in two steps. First, given the requirements of the passenger, (the destination is the most important requirement), the central authority is the *Level i Controller* that involves both origin and destination. This *Level i Controller* is committed to solving the shortest path in a high level problem (each node is a region). This path is only a first approach that guides us to find the real shortest path. Thus, *Level i Controllers* is executed first, then all *Level i-1 Controller* are executed concurrently and so on. Depending on the size of the journey, several hierarchical levels are necessary. Finally, the calculated route is sent to the *Segment Controllers* that are involved.
- Due to the dynamic structure of the problem, some parts of the system may change and new schedules must be calculated. The rescheduling is only calculated from the incidence to the destination. The management of backtrack-

ing is carried out in a way similar to the railway scheduling problem distributed by stations.

- As we have pointed out, the nature of the system makes the presence of a central authority necessary. However, due to the scalability of the system, the central authority has the same behaviour as a level controller. The central authority is the minimal level controller that involves both origin and destination. This level depends on the problem instance.

Conclusions

In the paper, we question the common assumption made in DisCSP literature in which each agent has just a single variable. Many real problems can be modelled as a CSP, but they cannot be solved by using DisCSP techniques due to the exploitation in message passing. Thus, new distributed techniques must be developed to solve large instances of real problems. In this paper, we present a general distributed model for solving large-scale problems and some guidelines for distributing these problems by relaxing the above assumption. We present two real-life problems which can be modelled as a distributed problem. They manage several variables per agent in order to solve these problems in a reasonable time. These problems follow some of the presented guidelines.

References

- Abril, M.; Barber, F.; Ingolotti, L.; Salido, M.; Tormos, P.; and Lova, A. 2007. An assessment of railway capacity. *Transportation Research Part E-Logistics and Transportation Review*, to appear.
- Abril, M.; Salido, M.; and Barber, F. 2007. DFS-tree based heuristic search. In *Proceeding of the Seventh Symposium on Abstraction, Reformulation and Abstraction (SARA'07)*, LNAI 4612, to appear.
- Bacchus, F., and van Beek, P. 1998. On the conversion between non-binary and binary constraint satisfaction problems. In *proceeding of AAAI-98* 311–318.
- Cordeau, J.; Toth, P.; and Vigo, D. 1998. A survey of optimization models for train routing and scheduling. *Transportation Science* 32:380–446.
- Ezzahir, R., B. C. B. M., and Bouyakhf, E.-H. 2007. Dischoco: A platform for distributed constraint programming. In *Proceedings of IJCAI-07 Workshop on Distributed Constraint Reasoning*.
- Faltings, B., and Yokoo, M. 2005. Introduction: Special issue on distributed constraint satisfaction. *Artificial Intelligence* 161:1–5.
- Giret, A., and Botti, V. 2004. Holons and Agents. *Journal of Intelligent Manufacturing* 15:645–659.
- HMS, P. R. 1994. *HMS Requirements*. <http://hms.ifw.uni-hannover.de/>: HMS Server.
- Koestler, A. 1971. *The Ghost in the Machine*. Arkana Books.
- Salido, M., and Barber, F. 2006. Distributed CSPs by graph partitioning. *Applied Mathematics and Computation*. Ed. Elsevier Science 183:491–498.
- Salido, M.; Abril, M.; Barber, F.; Ingolotti, L.; Tormos, P.; and Lova, A. 2007. Domain-dependent distributed models for railway scheduling. *Knowledge Based Systems*. Ed. Elsevier Science 20:186–194.
- Salido, M.; Giret, A.; and Barber, F. 2003. Distributing Constraints by Sampling in Non-Binary CSPs. In *IJCAI Workshop on Distributing Constraint Reasoning* 79–87.
- Salido, M. 2007. Distributed CSPs: Why it is assumed a variable per agent? In *Proceeding of the Seventh Symposium on Abstraction, Reformulation and Abstraction (SARA'07)*, LNAI 4612, 407–408.
- 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.
- Silaghi, M., and Faltings, B. 2005. Asynchronous aggregation and consistency in distributed constraint satisfaction. *Artificial Intelligence* 161:25–53.
- 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.
- Yokoo, M.; Durfee, E.; Ishida, T.; and Kuwabara, K. 1998. Distributed constraint satisfaction algorithm for complex local problems. *Third International Conference on Multiagent Systems (ICMAS-98)* 372–379.
- Zelinkovskyn, R. Global automated transport system. <http://www.global-transportation.com>.