Flexible Integration of Planning and Information Gathering

David Camacho and **Daniel Borrajo** and **José M. Molina** and **Ricardo Aler**

Universidad Carlos III de Madrid, Computer Science Department

Avenida de la Universidad n 30

CP 28911, Leganés, Madrid, Spain

{dcamacho, dborrajo, molina}@ia.uc3m.es, aler@inf.uc3m.es

Abstract

The evolution of the electronic sources connected through wide area networks like Internet has encouraged the development of new information gathering techniques that go beyond traditional information retrieval and WEB search methods. They use advanced techniques, like planning or constraint programming, to integrate and reason about hetereogeneous information sources. In this paper we describe MAPWEB. MAPWEB is a multiagent framework that integrates planning agents and WEB information retrieval agents. The goal of this framework is to deal with problems that require planning with information to be gathered from the WEB. MAPWEB decouples planning from information gathering, by splitting a planning problem into two parts: solving an abstract problem and validating and completing the abstract solutions by means of information gathering. This decoupling allows also to address an important aspect of information gathering: the WEB is a dynamic medium and more and more companies make their information available in the WEB everyday. The MAPWEB framework can be adapted quickly to these changes by just modifying the planning domain and adding the required information gathering agents. For instance, in a travel assistant domain, if taxi companies begin to offer WEB information, it would only be necessary to add new planning operators related to traveling by taxi, for a more complete travel domain. This paper describes the MAPWEB planning process, focusing on the aforementioned flexibility aspect[1].

Introduction

In recent years there has been a lot of work in Web information gathering. Information gathering intends to integrate a set of different information sources with the aim of querying them as if they were a single information source (Fan and et al. 1999; Lambrecht and Kambhampati 1997). Many different kinds of systems, named *mediators*, have been developed. They try to integrate information from multiple distributed and heterogeneous information sources, like database systems, knowledge bases, web servers, electronic repositories...(an example is the *SIMS* (Knoblock et al. 1997) architecture). In order that these systems are

[1]Current affiliation of David Camacho: Universidad Autonoma de Madrid, Computer Science Department, email: david.camacho@uam.es

practical, they must be able to optimize the query process by selecting the most appropriate WEB sources and ordering the queries. For this purpose, different algorithms and paradigms have been developed. For instance, *Planning by Rewriting* (*PbR*) (Ambite and Knoblock 1998; 1997) builds queries by using planning techniques.

Other examples of information gathering systems are:

- Ariadne (Knoblock et al. 1997): This system includes a set of tools to construct wrappers that make WEB sources look like databases. It also uses mediation techniques based on SIMS.

- Heracles (Knoblock et al. 2001): This framework is used to develop different information assistant systems that use a set of information agents (Ariadne, Theseus, Electric Elves). It uses a dynamic, hierarchical constraint propagation network to integrate the different information sources. Two assistant systems have been implemented: *The Travel Planning Assistant* (specialized in assisting tourists to plan their trips) and *The WorldInfo Assistant* (for a user-specified location, it retrieves information like weather, news, holidays, maps, airports, ...).

- WebPlan (Hullen, Bergmann, and Weberskirch 1999): it is a WEB assistant for domain-specific search on the Internet based on dynamic planning and plan execution techniques. The existing planning system *CAPlan* has been extended in different ways in order to deal with incomplete information, information seeking operators, user interaction, and interleaving planning and execution. WebPlan is specialized in localizing specific PC software on the Internet.

Some of the previous approaches use planning techniques to select the appropriate WEB sources and order the queries to answer generic user queries. That is, they use planning as a tool for selecting and sequencing the queries. In this paper we describe MAPWEB, an information gathering system that also uses planning, but with a different purpose (some preliminary work can be found in (Camacho, Molina, and Borrajo 2000; Camacho et al. 2001)). MAPWEB uses planning for both determining the appropriate generic sources to query and solving actual planning problems. For instance, in this paper, the MAPWEB framework is applied

to a travel planning assistant domain (e-tourism),[2] where the user needs to find a plan to travel along several places. Each plan not only determines what steps the user must perform, but which information sources should be accessed. For instance, if a step is to go from A to B by plane, the system gives the user the information of what airplane companies should be consulted for further information.

This domain is similar to the travel planning assistant built using the Heracles framework. However, Heracles constrained network, which is a kind of plan schema, needs to be reprogrammed everytime the planning domain changes. MAPWEB tries to be more flexible by using planning techniques to create the plans. For instance, if it is desired to add a new information source to the system, it is only necessary to change the planning domain instead of reprogramming the plan schema by hand. For instance, if taxi fares were made suddenly available in the WEB, it would only be necessary to add a move-by-taxi operator among with the associated WebAgent.[3] Actually, MAPWEB can handle planning operators which are not associated to any information source (because it is not yet available). This is useful, because even if no specific information is supplied, at least the user is told that he can fulfill that step by taxi.

MAPWEB **System Architecture**

MAPWEB is structured into several logic layers whose purpose is to isolate the user from the details of problem solving and WEB access. More specifically, we considered four layers between users and the WEB : *the physical world* (the users), *the reasoning layer* (that includes user agents, planning agents, and control agents), *the access information layer* (that contains WebAgents to retrieve the desired information), and *the information world* (which represents the available information). This four-layer architecture can be seen in Figure 1.
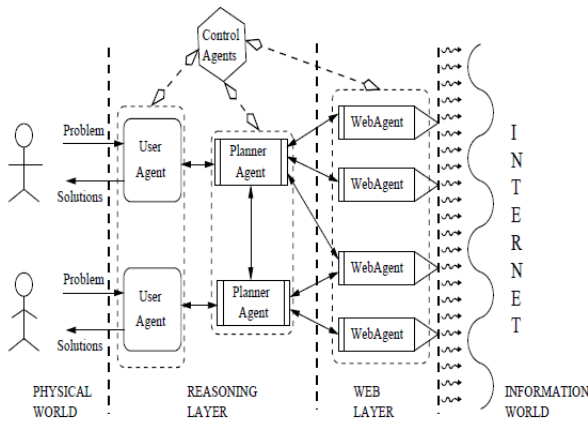

Figure 1: MAPWEB three-layer architecture.

---

[2]This domain is a modified version of the Logistics domain.

[3]A WebAgent is an information agent specialized in consulting a particular information source.

MAPWEB deploys this architecture using a set of heterogeneous agents.Next, each of these types of agents will be described:

- **UserAgents**: They pay attention to user queries and display to users the solution(s) found by the system. When an UserAgent receives problem queries from the users, it gives it to the PlannerAgent and when it answers back with the plans, they provide them to the user.

- **ControlAgents**: They handle several control functions like the insertion and deletion of agents in the system and communication management.

- **PlannerAgents**: They receive an user query, build an abstract representation of it, and solve it by means of planning. Then, the PlannerAgent fills in the information details by querying the WebAgents. The planner that has been used by the PlannerAgent is PRODIGY4.0 (Veloso et al. 1995).

- **WebAgents**: Their main goal is to fill in the details of the abstract plans obtained by the PlannerAgents. They obtain that information from the WEB.

The way these agents cooperate is as follows. First, the user interacts with the UserAgent to input his/her query. The query captures information like the departure and return dates and cities, one way or return trip, maximum number of transfers, and some preference criteria. This information is sent to the PlannerAgent, which transforms it into a planning problem. This planning problem retains only those parts that are essential for the planning process, which is named the *abstract representation* of the user query. PRODIGY4.0 provides several abstract solutions to the user query. The planning operators in the abstract solutions require to be completed with actual information that can be retrieved from the WEB. To accomplish this, the PlannerAgent sends information queries to specialized WebAgents, that return several records for every information query. Then, the PlannerAgent integrates and validates the solutions and returns the data to the UserAgent, which in turn displays it to the user. MAPWEB agents use a subset of the KQML speech acts (Finin et al. 1994). The whole process will be described in full detail in the next section.

The Planning Process

As mentioned before, in MAPWEB, the information gathering process is carried out by a set of WebAgents, but this process is guided by the PlannerAgent that reasons about the requested problem and the different information sources that would be available.

The planning process is divided into two parts: solving an abstract problem, and completing it with information gathered from the WEB. Planning is decoupled this way because of two reasons:

- The abstract planning problem is easier to solve by classical planners. This is because if all the information about all the available flights, all possible trains, etc. was included in the planning process, planning would be unfeasible.

- It is not necessary to access the WEB during the planning process. Queries to the WebAgents are carried out only when plans are ready. This allows to reduce the number of queries, because only those queries that are required by the solution are ever made.

Planning works as follows. First, the PlannerAgent receives a query from any UserAgent. This query is analyzed and translated into an abstract planning problem. Second, the PlannerAgent uses its own skills and knowledge about the problem and tries to solve it. If the solving process is successful, the PlannerAgent has a set of abstract solutions. These solutions are too general and only have the essential information for the planning process. Those abstract solutions need specific information to complete and validate them. The PlannerAgent builds a set of information queries (queries to other agents in the system to request specific information). It is important to try to optimize the number of queries due to the high number of possible instantiations. When the queries have been built, the PlannerAgent selects the set of WebAgents that will be asked. Finally, when the WebAgents answer with the information found in the WEB (if WebAgents are successful) the PlannerAgent integrates all the specific information with the abstract solutions to generate the final solutions that will be sent to the UserAgent. In Figure 2 the modular description of the planning process is shown.


Figure 2: Planning Process developed by the PlannerAgent. The user query is transformed into an abstract planning problem, which is subsequently solved by PRODIGY4.0. Each solution is partially instantiated by means of domain dependent heuristics. Every operator in a solution generates several WEB queries, which is sent to the appropriate WebAgents by using the agent hierarchy. The agents return several records, that are used to complete and validate the abstract solutions.

The next subsections explain this process in detail by focusing in the data structures used by each of the relevant agents: the user query generated by the UserAgent, the abstract problem, the abstract solutions, the specific knowledge used by the PlannerAgent, and finally the specific information records retrieved by the WebAgents.

The User Query

The planning process starts when the user supplies a problem to be solved. A user query is a sequence of stages. Each stage is a template that represents a leg of the trip, and contains several fields to be filled by the user. Table 1 shows an

Leg	Stage	Date	Restrictions	N Transfers
1	Turin → Madrid	Sep. 11th	Plane or train	0 or 1
2	3 nights stay	Sep. 11th	< 15.000 pts	-
3	Madrid → Toledo	Sep. 14th	Plane or train	0 or 1
4	Toledo → Turin	Sep. 14th	Plane or train	0 or 1

Table 1: A user problem to go from Turin to Madrid by airplane or train.

instance of a possible user query. It will be used to illustrate the rest of the article. This query is then sent to the PlannerAgent.

The planning domain and the abstract solutions

In order to solve abstract problems, PRODIGY4.0 requires a domain where the planning operators are described. Using planning at this stage (instead of using pre-programmed plans) provides two main advantages:

1. *Flexibility:* the system can be adapted to many different versions of travel domains and problems by just changing the domain description or the abstract problem generation method, respectively.

2. *Easy integration of new* WEB *sources.* The WEB is a dynamic medium: more and more companies make their information available in the web everyday. If a new information source (like taxi fares) is made available, MAP-WEB can be adapted quickly by just adding a new planning operator and establishing a relation with a WebAgent specialized in gathering the information from the WEB.

The abstract problem would be given to the PlannerAgent planner (PRODIGY4.0) which would obtain several possible abstract solutions. In this case, the planner would reply with the plans shown in Figures 3 (solutions with 0 transfers) and 4 (1 transfer).

```
Problem: 0-Transfers
Solution 1:
 <travel-by-airplane user1 plane0 airport0 airport2>
 <move-by-local-transport user1 lbus2 bustop20 trainstat21 city2>

Solution 2:
 <move-by-local-transport user1 lbus0 bustop00 trainstat01 city0>
 <travel-by-train user1 train0 trainstat0 trainstat2>
```

Figure 3: Abstract solutions generated by PRODIGY4.0 for Leg 1 with 0-Transfers.

This is a set of abstract plans that contain no actual details. Some of the plan steps might not even be possible because, for instance, there are no companies linking two cities. Therefore, those plans need to be *validated* and *completed*. The PlannerAgent acomplishes this task in the following way:

1. The abstract steps in the solution contain unbound variables that relate to transfer cities. They need to be bound

```
Problem: 1-Transfers
Solution 1:
 <travel-by-airplane user1 plane0 airport0 airport1>
 <move-by-local-transport lbus1 bustop10 trainstat11 city1>
 <travel-by-train user1 train1 trainstat1 trainstat2>

Solution 2:
 <travel-by-airplane user1 plane0 airport0 airport1>
 <travel-by-airplane user1 plane1 airport1 airport2>
 <move-by-local-transport lbus2 bustop20 trainstat20 city2>
..............
```

Figure 4: Abstract solutions generated by PRODIGY4.0 for Leg 1 with 1-Transfers.

```
Agent
   |___UserAgent0
   |---ReasonerAgent
                  |___PlannerAgent0
   |___WebAgents
            |___Travel
                     |___Fly: WebAgent-Iberia,WebAgent-Amadeus-Flight
                     |___Train: WebAgent-Renfe,WebAgent-RailEurope
                     |___Bus
            |___Hotel: WebAgent-Amadeus-Hotel
            |___Car
   |___ControlAgents
                  |___ManagerAgent0
                  |___CoachAgent0
```

Figure 5: Agents Hierarchy. It describes all the available agents in MAPWEB and their information gathering skills.

before the WebAgents are queried. The PlannerAgent restricts the number of bindings by applying a *geographic heuristic*. This is achieved as follows:

- If the origin and arrival cities belong to the same country, only the cities in that country are considered as possible transfer cities.
- Else, if the origin and arrival cities belong to the same continent, only the cities of that continent are considered.
- Otherwise, all cities are considered.

In the case of the first leg of the trip, as Turin and Madrid belong to Europe, we extract the cities that belong to this continent (currently, about 30).

Table 2 displays the queries that would be generated in this case.

Query send to the WebAgents	N? Transfers
(travel-by-airplane user1 plane0? Turin Toledo)	0
(travel-by-train user1 train0? Turin Toledo)	0
(travel-by-airplane user1 plane0? Turin Alicante)	1
(travel-by-airplane user1 plane0? Turin Barcelona)	1
(travel-by-airplane user1 plane0? Turin Paris)	1
(travel-by-train user1 train0? Turin Madrid)	1
…	…

Table 2: Queries partially instantiated

2. Planning operators of the abstract solutions and web sources are related by means of a WebAgent *hierarchy*. This *hierarchy* is used by the PlannerAgent to select the relevant WebAgents that will be used to obtain the information. This hierarchy allows the PlannerAgent to know which WebAgents know how to retrieve the required information. In Figure 5 a description of this hierarchy is shown.

3. Finally the PlannerAgent uses the previous information to build a set of queries that will sent to the selected WebAgents. If a planning operator is repeated in different abstract solutions, it is only considered once, to avoid repeating queries.
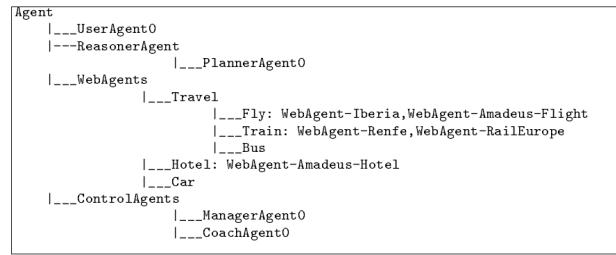
For instance, in the solutions for 1-Transfer problems the operator (`<travel-by-airplane user1 plane0 airport0 airport1>`) would be translated as shown in Table 3:

Query send to the WebAgents	WebAgent
(travel-by-airplane user1 plane0? Turin Toledo)	Iberia, Amadeus-Flights
(travel-by-train user1 train0? Turin Toledo)	Renfe, RailEurope
(travel-by-airplane user1 plane0? Turin Alicante)	Iberia, Amadeus-Flights
(travel-by-airplane user1 plane0? Turin Barcelona)	Iberia, Amadeus-Flights1
(travel-by-airplane user1 plane0? Turin Paris)	Iberia, Amadeus-Flights
(travel-by-train user1 train0? Turin Madrid)	Renfe, RailEurope
…	…

Table 3: Queries partially instantiated to the appropriate WebAgents.

Those queries (and all the additional information given by the UserAgent) are sent to several WebAgents that know about airplane travel, so that variable `plane0?` is instantiated as well.

Filling the Abstract Solutions

The information queries are sent to the selected WebAgents with the specific data (departure and arrival times, travel cost, etc...) and the query that the PlannerAgent needs. With this information the WebAgents automatically build the specific WEB query that will be sent to the WEB information sources the agent is specialized in. For every query, each WebAgent will return to the PlannerAgent a list of *records* by filling a template whose structure is shared by all the agents (there exist different templates depending on the kind of information). In Table 4 some of the retrieved flight-records and train-records provided by different WebAgents are shown for the Leg 1 in the travel example.

Finally, those records are received by the PlannerAgent that will use them to complete the abstract solutions. If the WebAgents return no records for a step of the abstract solution, that particular solution is rejected. However, it is important to remark that if it is known in advance that there are no WEB sources to complete a particular step (for instance, `<MOVE-BY-LOCAL-TRANSPORT taxi ...>`), then the user is told that s/he has to carry out that step, even though no specific information about that step

13

Inf-FLIGHTS	record1	record2	record3	Inf-TRAINS	record1	record2	record3
WebAgent	**Iberia**	**Amadeus**	**Amadeus**	**WebAgent**	**Renfe**	**Renfe**	**Renfe**
air-company	Iberia	Iberia	Portugalia	train-company	RENFE	RENFE	RENFE
http-address	w3.iberia.es	null	null	http-address	w3.renfe.es	w3.renfe.es	w3.renfe.es
flight-id	IB8797	IB8819	NI711	train-id	07054	07056	07058
ticket-fare	70641	null	null	ticket-fare	780	780	780
currency	ESP	ESP	ESP	currency	ESP	ESP	ESP
flight-duration	3h45min	2h00min	2h10min	departure-city	MAD	MAD	MAD
airp-depart-city	TRN	TRN	TRN	departure-date	11-09-01	11-09-01	11-09-01
departure-date	11-09-01	11-09-01	11-09-01	departure-time	6:30	8:30	10:14
airp-arrival-city	MAD	MAD	MAD	arrival-city	TOL	TOL	TOL
return-date	null	null	null	arrival-date	11-09-01	11-09-01	11-09-01
class	Tourist	null	null	arrival-time	7:53	9:47	11:30
n? passengers	1	1	1	class	Tourist	Tourist	Tourist
round-trip	one-way	one-way	one-way				

Table 4: Retrieved records by the WebAgents.

is attached. The set of completed solutions are finally sent to the UserAgent that requested the information.

Experimental Evaluation

The aim of this section is to carry out several experiments with MAPWEB to evaluate its performance. First, the example-trip we have used to illustrate the previous sections will be tested. Second, a set of problems given by the user will be evaluated to analyze the average behaviour of the system.

Table 5 summarizes the example-trip (Turin to Toledo and back). To solve this problem, a team of nine agents was used. They include all the agents displayed in the agents-hierarchy of Figure 5. In particular, airplane, train, and hotel WebAgents have been used. The next parameters have been measured:

- Validated abstract solutions/abstract solutions ratio (*val.sols/abs.sols*). This value measures how many abstract solutions provided by the planner were validated by the information provided by the information gathering agents.

- Number of instantiated solutions. It shows all the possible solutions to the user problem. The solutions are computed using the gathered records. The PlannerAgent uses the k validated abstract solutions that contain l_i abstract operators. If there are b_{ij} retrieved records for the j-th operator of the i-th solution, then the number of possible instantiated solutions is:

$$\text{Number of solutions} = \sum_{i=1}^{k} \prod_{j=1}^{l_i} b_{ij}$$

- Number of WEB Queries. This represent all the queries made by the WebAgents to retrieve the specific information.

- Number of gathered records (duplicated records are removed).

- Time. It includes planning time and WEB gathering time.

Everyone of the previous parameters is measured for both 0 and 1 transfers (0-T and 1-T). In this example, there are no solutions for the 0 transfers because it is impossible to complete the fourth leg of the trip (there is no way to go from Toledo to Turin directly). On the other hand, there are thousands of possible combinations when 1 transfer is allowed. It is important to remark that even though when 1 transfer is used, it takes several thousand seconds to find the solutions, only 1 second per leg is spent for actual planning.

We have also tested a set of 38 problems with different configurations of MAPWEB. The problems include 15 trips within Spain, 15 within Europe, and 8 Intercontinental ones. Each problem has been tried with 0 and 1 transfers. The results are shown in Table 6. This experiment shows in practice the flexibility of MAPWEB when it is necessary to add new information sources. In order to add a new information source to MAPWEB, it is only necessary to add a planning operator and the related WEB agent. The configurations that have been used are as follows:

- *N0*: only one WebAgent specialized in retrieving information from a particular Airplane Company (Iberia Airlines[4]) was considered.

- *N1*: different WebAgents specialized in gathering information of the same kind (flight information) were used: WebAgent-Iberia, WebAgent-Avianca, WebAgent-Amadeus-Flights, WebAgent-4Airlines-Flights. The two last ones are meta-searchers engines.

- *N2*: only two WebAgents specialized in gathering information of the same type (train information) were used: WebAgent-Renfe, WebAgent-RailEurope.

- *N3*: finally this configuration integrates all the previous WebAgents, that is, agents for retrieving both flight and train information (N3=N1+N2).

In Table 6, we observe the following:

- With respect to N0, as it could be expected, many more solutions are found when 1 transfer legs are allowed (999.3 vs. 7.1). It can also be observed that MAPWEB

---

[4]www.iberia.com/iberia_es/home.jsp

Leg	Stage	Val. sols per abs. sols		Number of solutions		Number of queries		Number of records		Time (seconds)	
		0-T	1-T	0-T	1-T	0-T	1-T	0-T	1-T	0-T	1-T
1	Turin → Madrid	0.5	0.667	2	1829	4	43	20	135	112.465	962.938
2	3 nights stay	1	1	6	6	1	1	20	20	62.384	62.384
3	Madrid → Toledo	0.5	0.333	12	797	4	43	22	92	75.030	1692.839
4	Toledo → Turin	0	0.333	0	432	4	43	0	67	76.236	3874.948

Table 5: MAPWEB request for the example-trip, with 0 and 1 transfers.

Config.	Number of solutions		Solved problems		Number of queries		Time (seconds)	
	0-T	1-T	0-T	1-T	0-T	1-T	0-T	1-T
N0	7.1	999.3	65.7%	74.3%	1	26.9	65.6	1485.7
	$\sigma = 6.7$	1480.5					$\sigma = 10.2$	1319.2
N1	10.9	1338.1	94.2%	97.1%	4	91.2	162.4	2243.6
	$\sigma = 8.0$	1725.8					$\sigma = 200.5$	1321.8
N2	3.7	3.7	25.7%	40.0%	2	51.4	70.1	1314.4
	$\sigma = 7.9$	7.9					$\sigma = 23.0$	1124.3
N3	12.5	1340.2	94.3%	97.1%	6	143.9	165.3	2666.6
	$\sigma = 8.8$	1724.4					$\sigma = 199.6$	1298.8

Table 6: Summary of the results for 38 user problems, with 0 and 1 transfers (0-T and 1-T).

cannot find a solution for some problems, although the number of problems solved increases for the 1-T option (74.3% vs. 65.7%). However, the number of queries and the time required to fulfill them also increases quickly. It is also noticeable that standard deviations are rather large. This is because user problems can be very different; some of them can be solved quickly because there are few retrieved records, whereas other problems can have many possible solutions.

- N1 enlarges N0 by including more airplane companies. MAPWEB does not find many more solutions per problem, because most of the user problems are within Europe, where Iberia (the only agent in N0) offers many flights. However, many more problems are solved (94.2% vs. 65.7% with 0 transfer, and 97.1% vs. 65.7%). Although the number of queries is multiplied by 4 in N1, the time required to fulfill them has been only doubled (162.4 vs. 65.6 for 0-T and 2243.6 vs. 1485.7 for 1-T). Time is doubled because even though the four WebAgents work in parallel, all the retrieved records must be analyzed by a single PlannerAgent.

- N2 displays the results when only train travels are allowed. Only a few problems can be solved: 25.7% with 0-T and 40.0% with 1-T, and very few solutions per problem are found (3.7). This is clearly due to the smaller number of possibilities of fullfilling travels using only trains vs. using airplanes.

- N3 integrates both airplane and train companies. Compared to N1, almost the same number of user problems are solved (94.3% vs. 94.2% and 97.1% vs. 97.1%), although some more solutions per problem are found (12.5 vs. 10.9 and 1340.2 vs. 1338.1).

Conclusions and Future Work

The WEB is a dynamic medium: more and more companies make their information available in the web everyday. WEB information gathering systems need to be flexible to adapt to these rapid changes. In this paper we have described in detail MAPWEB, a multiagent framework that combines classical planning techniques and WEB information retrieval agents. MAPWEB decouples planning from information gathering, by splitting a planning problem into two parts: solving an abstract problem and validating and completing the abstract solutions by means of information gathering. Flexible information gathering is achieved by means of planning. In order to add a new information source to the system, only the planning domain has to be modified, besides adding the related WEB agent. For instance, in order to add information sources about train travels, MAPWEB planning domain has to be extended by adding a TRAVEL-BY-TRAIN operator and a set of WebAgents specialized in retrieving this kind of information.

In the future, several new skills will be developed for different agents in MAPWEB. These skills will try to improve the performance of the global system in two ways: by increasing the number and quality of solutions found by the agents, and by minimizing the time and computational resources used by MAPWEB to solve problems. We could summarize these main future lines as follows:

1. To develop Distributed Planning Skills. Cooperation between different PlannerAgents (desJardins and Wolverton 1999) would allow to distribute the computational cost of the plannings process. In the domain considered in this paper, it seems easy to develop distributed approaches because the problems considered have independent goals.

2. To develop Case-Based Planning Skills (CBP) (Bergmann and Wilke 1996; Veloso and Carbonell 1993). The Plan-

nerAgents would store old successful plans and use Case-Based Reasoning techniques, so that new planning problems can be solved by adapting previously solved plans which were similar. This would reduce enormously the planning process which might become computationally expensive.

3. Optimizing the query information heuristic. It should be possible to improve the heuristic used by the Planner-Agent adding more knowledge to the city-world structure, and therefore optimize the number of the information queries generated by those agents.

4. Reuse of information stored in WebAgents. These Agents can learn from experience, and reuse information retrieved previously to reduce WEB access (Howe and Dreilinger 1997; Lieberman 1995).

Acknowledgements

References

Ambite, J. L., and Knoblock, C. A. 1997. Planning by rewriting: Efficiently generating high-quality plans. In *In Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 706–713.

Ambite, J. L., and Knoblock, C. A. 1998. Flexible and scalable query planning in distributed and heterogeneous environments. In *In Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems*, 3–10.

Bergmann, R., and Wilke, W. 1996. Paris: Flexible plan adaptation by abstraction and refinement. In *Proceedings of the Workshop on Adaptation in Case-Based Reasoning, th European Conference on Artificial Intelligence*. John Wiley & Sons.

Camacho, D.; Molina, J. M.; Borrajo, D.; Aler, R.; and Madrid, I. 2001. Mapweb: Cooperation between planning agents and web agents. *Information & Security: An International Journal. Volume 7* 7:209–238.

Camacho, D.; Molina, J. M.; and Borrajo, D. 2000. A multiagent approach for electronic travel planning. In *In Proceedings of the Second International Bi-Conference Workshop on Agent-Oriented Information Systems (AOIS-2000*, 200–0. AAAI Press.

desJardins, M., and Wolverton, M. 1999. Coordinating a distributed planning system. *AI Magazine* 20(4):45–53.

Fan, Y., and et al. 1999. Adaptive agents for information gathering from multiple, distributed information sources.

Finin, T.; Fritzson, R.; McKay, D.; and McEntire, R. 1994. Kqml as an agent communication language. In *Proceedings of the third international conference on Information and knowledge management*, CIKM '94, 456–463. New York, NY, USA: ACM.

Howe, A. E., and Dreilinger, D. 1997. Savvysearch: A meta-search engine that learns which search engines to query. *AI Magazine* 18.

Hullen, J.; Bergmann, R.; and Weberskirch, F. 1999. Webplan: Dynamic planning for domain-specific search in the internet. In *in the internet, in Workshop Planen und Konfigurieren (PuK-99*.

Knoblock, C. A.; Minton, S.; Ambite, J. L.; Ashish, N.; Modi, P. J.; Muslea, I.; Philpot, A. G.; and Tejada, S. 1997. Modeling web sources for information integration.

Knoblock, C. A.; Minton, S.; Ambite, J. L.; Muslea, M.; Oh, J.; and Frank, M. 2001. Mixed-initiative, multi-source information assistants. In *In The Tenth International World Wide Web Conference (WWW10), Hong Kong*, 697–707. ACM Press.

Lambrecht, E., and Kambhampati, S. 1997. Planning for information gathering: A tutorial survey. Technical report, ASU CSE Techinal Report.

Lieberman, H. 1995. Letizia: An agent that assists web browsing. In *IJCAI (1)*, 924–929.

Veloso, M. M., and Carbonell, J. G. 1993. Derivational analogy in prodigy: Automating case acquisition. In *Storage, and Utilization. Machine Learning*.

Veloso, M.; Carbonell, J.; Prez, A.; Borrajo, D.; and Blythe, J. 1995. Integrating planning and learning: The prodigy architecture. *Journal of Experimental and Theoretical Artificial Intelligence* 7:81–120.