

Intelligent Agents in Portfolio Management

Katia Sycara

Keith Decker

Dajun Zeng

October 22, 1996

Abstract

The voluminous and readily available information on the Internet has given rise to exploration of Intelligent Agent technology for accessing, filtering, evaluating and integrating information. In contrast to most current research that has investigated single-agent approaches, we are developing a collection of multiple agents that team up on demand, depending on the user, the task and the situation, to access, filter and integrate information in support of user tasks. We are investigating techniques for developing distributed *adaptive* collections of information agents that coordinate to retrieve, filter and fuse information relevant to the user, task and situation, as well as anticipate user's information needs. Our approach is based on (1) *case-based task and situation models*, (2) *flexible organizational structuring*, and (3) *reusable agent architecture*. We are currently implementing the system in the domain of financial portfolio management.

1 Introduction

Due to advances in technology, diverse and voluminous information is becoming available to decision makers. This presents the potential for improved decision support, but poses challenges in terms of building tools to support users in accessing, filtering, evaluating and fusing information from heterogeneous information sources. Most reported research on Intelligent Information Agents to date has dealt with a user interacting with a single agent that has general knowledge and is capable of performing a variety of user delegated information finding tasks (e.g., [4]). For each information query, the agent is responsible for accessing different information sources and integrating the results. We believe that, given the current computational state of the art, a centralized agent approach has many limitations: (1) a single general agent would need an enormous amount of knowledge to be able to deal effectively with user information requests that cover a variety of tasks, (2) a centralized information agent constitutes a processing bottleneck and a "single point of failure", (3) unless the agent has beyond the state of the art learning capabilities, it would need considerable reprogramming to deal with the appearance of new agents and information sources in the environment, (4) because of the complexity of the information finding and filtering task, and the large amount of information, the required processing would overwhelm a single agent. Because of these reasons and because of the characteristics of the Internet environment, we employ a distributed collaborative collection of agents for information gathering.

We are currently working on a system where each user is associated with a set of agents which have access to the task and situation models and keep track of the current state of the task, situation, environment and user information needs. Based on this knowledge, the agents decide what information is needed and initiate collaborative searches with other agents to get the information. During search, the agents communicate with each other to request or provide information, find information sources, filter or integrate information, and negotiate to resolve conflicts in information and task models. The returned information is communicated to display agent or agents that possibly combine it with information from other sources (e.g. the user) and/or filter it for appropriate display to the user.

This paper focuses on the design of such a system of agents for the task environment of financial portfolio management, and on the key issues that we will be addressing. These issues include:

- **Gathering and integrating diverse information sources with collaborating software agents.** Because of the volume, complexity, and dynamic nature of the information available to support a user's goals, it is impossible for a single computational agent to find, access, and integrate all that information in a timely manner.
- **Case-Based User, Task and Situation Models.** The utility of information depends on the user, task and situational context. We propose to model user, task, and situations in cases. Cases organize information gathering plan fragments under currently valid user, task, and situation features; focus distributed agent information searches; and guide information filtering and integration. Case updates resulting from agent information gathering activities refine these models and reflect the current situation.
- **Adaptive integration of planning, coordination, scheduling, and execution.** We will explore the construction of individual agents that have the ability to be highly autonomous, quickly adapt to changes in the current environmental situation, and yet still be socially situated, balancing predictability and responsiveness.

In our system, Case-based reasoning provides meta-level control and activation of agents. Depending on the task, user and situation, case based retrieval selects current planning goals, information needs and information gathering goals. Based on the plans and information gathering goals, agent teams are activated "on demand" to access, integrate and filter information to fulfill these goals. New information can be incorporated in the case base and may give rise to new plans and information gathering goals (and as a result activation of potentially different agent teams). The system has two types of agents, task agents and information agents. Task agents have information about tasks and associated information gathering goals. Information agents have models of information sources, information access strategies and associated task agents to whom the information should be returned. The reported work is a continuation of our previous work on multi agent information access, filtering and integration of everyday organizational tasks [17].

2 The Portfolio Management Domain

To evaluate our domain independent agent control, organization, coordination and architectural schemes, we have chosen financial portfolio management as a task domain. This is the task of providing an integrated financial picture for managing an investment portfolio over time, using the information resources already available over the Internet. This task environment has many interesting features, including:

- the enormous amount of continually changing, and generally unorganized, information available
- the variety of kinds of information that can and should be brought to bear on the task (market data, financial report data, technical models, analysts' reports, breaking news, etc.)
- the many sources of uncertainty and dynamic change in the environment
- information timeliness and criticality features that present the agents with hard and soft real-time deadlines for certain tasks
- resource and cost constraints—not all data are available for free
- relatively well-structured evaluation criteria and experimentally verifiable testbed where decisions supported by the system can be evaluated using real world data and feedback

The overall task in the portfolio management domain, as stated by modern portfolio theory [9], is to provide the best possible rate of return for a specified level of risk, or conversely, to achieve a specified rate of return with the lowest possible risk¹. Risk tolerance is one of the features that characterize the user of our system; other features include the user's investment goals (long-term retirement savings? saving for a house?) and the user's tax situation. In current practice, portfolio management is carried out by investment houses that employ teams of specialists for finding, filtering and evaluating relevant information. Current practice as well as software engineering considerations motivate our multi agent system architecture.

Previous work in the portfolio management domain (see [19] for one collection) has focused on the portfolio selection process (i.e., "stock picking") as opposed to portfolio monitoring—the ongoing, continuous, daily provision of an up-to-date financial picture of an existing portfolio. A multi-agent system approach is natural for portfolio monitoring because the multiple threads of control are a natural match for the distributed and ever-changing nature of the underlying sources of data and news that affect higher-level decision-making processes. A multi-agent system can more easily manage the detection and response to important time-critical information that could appear suddenly at any of a large number of different information sources. Finally, a multi-agent system provides a natural mapping of multiple types of expertise to be brought to bear during any portfolio management decision-making. Existing DAI techniques for resolving conflicting opinions, negotiation, and argumentation can be brought to bear on these problems [15, 16].

The overall portfolio management task has several component tasks. These include eliciting (or learning) user profile information, collecting information on the user's

¹Risk can be specified via a statement such as "I wish to be 95% certain that I endure no more than a 10% loss in value over any one year."

initial portfolio position, and suggesting and monitoring a re-allocation to meet the user’s current profile and goals. As time passes, assets in the portfolio will no longer meet the user’s needs (and these needs may also be changing as well). Our initial system focuses on this ongoing portfolio monitoring process.

The task of monitoring individual portfolio assets gives rise to a variety of concurrent goals, such as monitoring an asset currently being held (should we continue to hold it? sell some or all of it?), or monitoring the buying or selling of an asset. Buying and selling at this high level are not instantaneous transactions, but also require careful planning and monitoring of plan execution. For example, let us assume that in the process of monitoring the system comes to recommend that an asset be sold. The way in which it is sold will be determined in part by the reason for the sale—perhaps the asset is no longer performing as required for its role in the portfolios asset allocation mix. Perhaps it is performing *too* well, and there is a growing possibility that the asset has reached a peak. In this second case it is often prudent not to sell one’s entire holdings all at once, but to sell in phases and place the appropriate standing orders to protect the user from a sudden downturn (while avoiding worry over simple daily price fluctuations)². Thus the goal of selling an asset is not one that requires only a simple short sequence of actions, but one that requires careful planning and monitoring of that plan as it is executed, over an extended period of time.

3 Case Based Situation and Task Models

A real-world information gathering and decision making/aiding system, such as a portfolio management system, typically operates in a dynamically evolving environment, where information gathering can be user-, task- and situation dependent. Good models of the user, task, and situation help focus information gathering and filtering activities so that relevant information can be efficiently found. On the other hand, due to the unpredictability of the information and decision-making environment, new goals come into the system in the midst of execution, actions may fail due to exogenous events and need to be replanned, and there is incomplete information about the world.

We are proposing to model the information gathering task as a planning task where planning and execution are interleaved and search is guided by user, task and situation models. We believe that cases that incorporate user, task and situation models can effectively provide instantiations of the situation-dependent task structure and the associated team of information gathering agents. The case base contains cases of successful and unsuccessful information gathering episodes and information evaluation. After each information gathering cycle the case base is updated. Thus, learning is integrated with problem solving and is achieved automatically. This feature makes Case-Based reasoning very preferable in application domains with open and dynamically changing world model [18]. Case based reasoning offers three general advantages. First, since it relies on reusing specific experiences rather than reasoning from a general world model, it provides for more efficient planning. Second, since it can start with few cases in memory and incrementally acquire new cases based on a reasoner’s interactions with the world, it makes few assumptions about the completeness and correctness of world knowledge. Third, previous past failures warn the planner about the possibility

²There are of course tax considerations on sales as well.

of failure in current circumstances and, hence help avoid future failures.

Each case will be automatically labeled with identifying information such as user's name (a user could be a software agent), time of creation, time of modification, etc. So, a case will carry a complete audit trail of its origin and modifications. This information is essential for careful analysis of all the planning knowledge that is exchanged during system operation, and we will rely heavily on it during the initial stages of knowledge acquisition and development. In addition, this information can be used to evaluate the performance of the system under different planning scenarios, since each item will be clearly identified.

The library of cases we propose can be viewed as a case-based scheme for meta control that offers agents access to task and situation specific information. The agents use this information to focus search and filtering. The asynchronous information gathering activity of the agents results in new information about the world (new cases) that gets incorporated into the case base. Case base updates result in formation of new memory indices. These new indices, along with any new user inputs (e.g., information gathering requests, change in context) activate a new set of cases that reflect possibly new information seeking goals, and new sets of agents. In this sense, the case base can be viewed as (1) tracking the user's intentions and his/her evolving information needs, (2) reflecting changing situations, (3) recognizing new events, and (4) learning information retrieval tactics/heuristics.

Figure 1 shows an active case base which receives as input notification of events, either directly from the user or from new information that becomes available (e.g., from other software agents). Updates of the case-base with the addition of new information finding episodes and new indices are also considered an event that results in the activation of a new round of reasoning. In the following, we use an example to illustrate the general Case-based agent invocation process. Given a situation, the case-base calls a certain task assistant, say Analyst Tracking Agent, which the reasoning process determines are relevant to meeting the information gathering goals (which in turn might be information requests from Portfolio Manager Agent). Portfolio Manager Agent call upon Earnings analysis Agent, News Classifier, etc., to locate information from the infosphere directly (in the case of News Classifier) or indirectly (in the case of Earning Analysis Agent). Note that there is no hierarchy of agents here: in a different situation Earning Analysis Agent might have been called directly by the case-base reasoning process. After collaborating to find and filter information, Analyst Tracking Agent agent updates the case base with new information finding episodes that include a timestamp and the results of the search. If unexpected news ("unexpected" in terms of the current situation) has been found during information retrieval (such as major corporate merge news), the case-based process will interrupt the regular plan execution and take other emergency actions. For example, the case-base process might invoke the interface agent to notify the user right away.

From the perspectives of Case-Based reasoning, there are a number of important research questions that we need to address. The most fundamental question is *what constitutes a case?* A case in case memory describes a specific information scenario including: (1) information needs and goals which might come directly from the user or from other software agents, (2) global features which give an abstract characterization of the situation in which this information gathering operation takes place, (3) features of local nature which describe in detail the information about informa-

tion sources, inter-agent interactions, etc., (4) information retrieval plan skeleton, (5) feedbacks/evaluations with respect of acquired information and information retrieval effectiveness/efficiency, and (6) potential failures and re-trial, plan modification information.

At the most abstract level, we might index an asset monitoring task plan fragment using several broad situation characteristics. These characteristics are represented in the case base and used as indices for the retrieval of plan fragments, associated information gathering goals and associated lists of agents to be activated. Depending on the situation, at different times, different task and information plans could be active, thus activating different agent teams.

Goal-to-plan-for: We described some abstract monitoring goals above, such as monitoring a held position, monitoring an acquisition, or monitoring the sale of all or part of a position. This is a main index for plan fragment retrieval.

Asset type: U.S. stocks, U.S. government bonds, mutual funds, gold, etc. The high-level, abstract plans used for monitoring these different types of investments are themselves quite different.

Sector: With stocks, for example, this is the industry sector such as electronics or health care³. While this will only have a minimal effect on the high level abstract plan an agent uses, it will have a larger effect at more detailed levels. For example, the types of information (and the evaluation processes) that are needed to evaluate an investment in a financial institution such as a bank are very different from those need to evaluate the stock of a manufacturing company. It also offers an index for learning sector-wide information such as economic forecasts and cross company earnings reports effects (see Section 4.2).

Name: Again, the particular stock being monitored is an important index at the more detailed levels of plan retrieval, and will be used for indexing task execution times, storing historical data, and so on.

Ownership records: Information about how much was purchased, when, and at what price. This information interacts with the user's investment goals and tax situation to produce situation-specific buy/sell plans.

Tax status: Assets held in non-taxable trusts like an IRA require different treatment than those that are subject to taxes on realized capital gains.

Portfolio status: What else is in the portfolio? What is the current asset allocation, and what role does this asset play?

User profile: What are the user's investment goals for this portfolio? Expectations on return? Risk tolerance?

Other: Other information, such as brokerage commissions.

We are currently investigating other important issues such as efficient case indexing mechanism, case retrieval/matching approaches, and initial case collection⁴

³Actually, there are several well-defined hierarchical classification systems used in the industry.

⁴To address the initial case collection issues, we envision that a limited number of standard finance information gathering and decision making procedures (default cases) can serve well for the initial case-base. This is due to another nice property of Case-based reasoning: a partial case-base can be a starting point of a case-based system which will gather cases through user interaction and problem solving incrementally.

4 Organizational Structure

We propose a general system organization in which agents are directly activated based on the top-down elaboration of the current situation (as opposed to indirect activation via manager or matchmaker agents, or self-directed activation). These agent activations, guided by case-based retrieval according to the current situation, dynamically form an organizational structure that fits in with the user's current profile, tasks, and other situational features. This organization will change over time, but will also remain relatively static for extended periods (for example, while monitoring currently held investments during stable market periods). Information that is important for decision-making (and thus might cause an eventual change in organizational structuring) is monitored at the lowest levels of the organization and passed upward when necessary.

In this type of organization (see Figure 1), "task agents" or "task assistants" [17] continually interleave planning, scheduling, coordination, and the execution of domain-level problem-solving actions. Task agents interact with one another and with "information agents" or "information assistants" that encapsulate network information sources. Task agents retrieve, coordinate, and schedule plans based on local knowledge modulated by situational context. A task assistant decomposes an information request into information seeking goals and subgoals and interacts with the information assistants to gather the information. In this architecture, a task assistant does the final filtering and fusing of information before it passes it on to agents above it in the organizational structure (requesting agents). This incremental information fusion and conflict resolution increases efficiency and potential scalability (e.g., inconsistencies detected at the information-assistant level may be resolved at that level and not propagated to the task-assistant level) and robustness (e.g., whatever inconsistencies were not detected during information assistant interaction can be detected at the task-assistant level). In addition, a task assistant composes a new case that incorporates its findings to be stored in the case memory.

In this architecture, information-assistants would have models of their associated information sources, learn the reliability of those sources, as well as strategies for low-level information fusion and multiple methods for responding to information requests. As an example of the latter, a stock ticker monitoring agent might have several methods available to it that trade off time, cost, and quality:

- one or more sources of 15-minute delayed values (with varying reliabilities and average response delays)
- one or more sources of real-time quotations that charge a fee (more reliable response but still not guaranteed)
- the ability to guess a quote based on recent data and simple models (very fast but of low quality).

On the other hand, task-specific assistants have a model of the task domain, executable methods for performing the task, knowledge of an initial set of information-assistants relevant to their task and strategies for learning models of pertinent information-assistants.

Figure 1 shows a top-level portfolio manager agent which receives as input notification of events, either directly from the user or from the case base, or from information

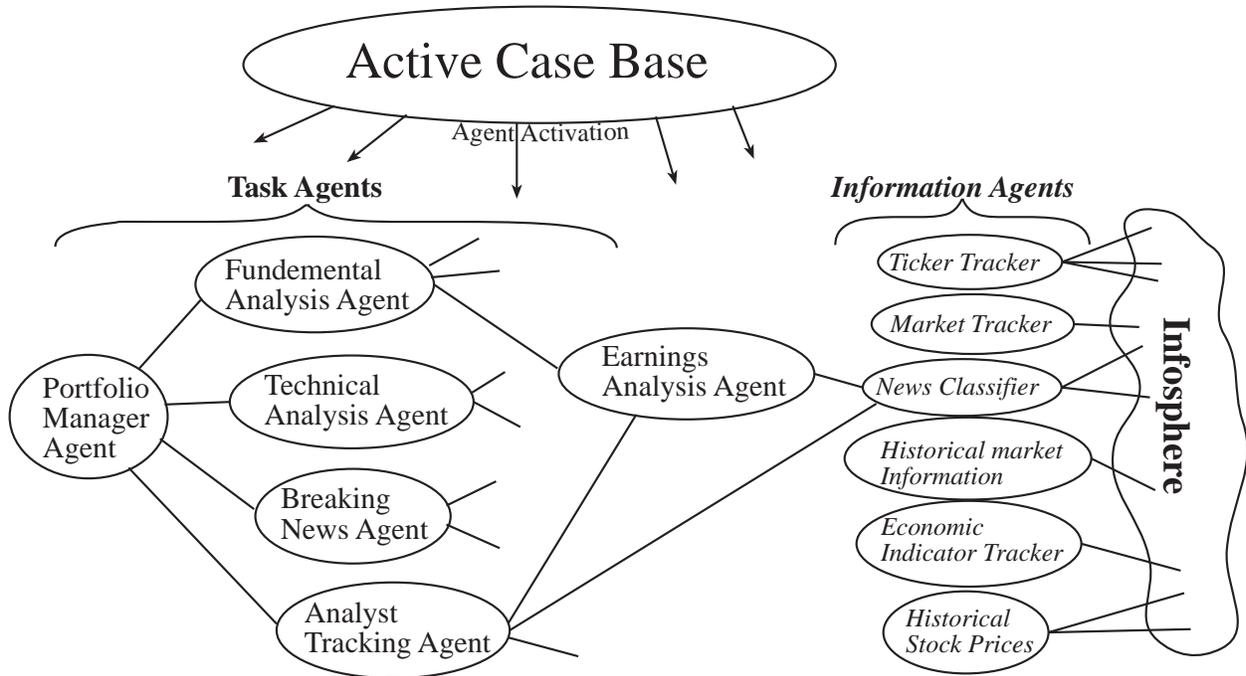


Figure 1: A “direct invocation” agent organization for a portfolio management system.

that becomes available (e.g., from task and information agents). Given the current situation, the portfolio manager agent:

- instantiates task plans and associated information gathering goals according to the current situation
- coordinates those plans with other agents (this includes task assignment actions that activate task assistants)
- schedules and monitors the execution of its local actions.

In Figure 1, the fundamental, technical, news, and outside-analyst task agents have been activated in this manner. These agents are task assistants that can either locate information via information assistants, or by calling upon other task agents. There is not a strict hierarchy of agents—the same task and information agents may be called upon by different parts of the portfolio management organization. After collaborating to find and filter information, task agents update the case base with new information finding episodes that include a time stamp and the results of the search.

This architecture has potential advantages and drawbacks. The advantages include:

- There is a finite number of task assistants that each agent communicates with.
- Because information processing is done by all the task agents at every level (rather than by having one task agent receive all data from every information agent) we avoid having a single computational bottleneck point.
- The task assistants are responsible for checking information quality, filtering irrelevant information, recognizing important information, and integrating information from heterogeneous information sources for their respective tasks.

- The task assistants are responsible for activating relevant information assistants and coordinating the information finding and filtering activity for their task.

All the above characteristics, by imposing some structure through definition of task assistants, contribute to overall system responsiveness. On the other hand, there are potential drawbacks:

- The portfolio manager is a single point of failure. Such failures can be mitigated by expending the resources to have, for example, a redundant portfolio manager that takes over in case of failure.
- Each task assistant also constitutes a “single point of failure for that task”. This can be mitigated by having more than one task assistant (either clones of each other or not). In the case where two different assistants exist for the same task, the task assistants must negotiate to resolve inconsistencies. We propose to explore the use of negotiation strategies for resolving inconsistencies.

In the stock portfolio example, task assistants for areas such as earnings analysis might be replicated and allowed to specialize on various industry groups (one agent to handle banking industry earnings, one for manufacturing, etc.). Such agents might begin as clones, but learn specialized case information. In the event of a failure, a non specialist would still be able to retrieve useful plans for a task inside its area of expertise, but outside of its specialty.

4.1 The Portfolio Monitoring Task

We can represent the plans that are retrieved using TÆMS task structures [2, 1]. TÆMS task structures are based on abstraction hierarchies, where task plans are elaborated via a “subtask” relationship into acyclic directed graphs that have actions, called executable methods, at their leaves. Such structures are compatible with most planning representations, and provide the necessary information both for scheduling activities that arise from multiple plans [6], and for coordinating the activities of multiple agents [3]. As shown in Figure 2, a top level portfolio management agent interacts graphically and textually with the user to acquire information about the user’s profile and goals; as mentioned earlier, we will assume in this paper that the system has gone through an initial usage period and has reached a “steady state” of monitoring the current portfolio.

Such a monitoring task includes gathering opinions from various task experts, integrating this information, and then making or updating the recommendation (such as buy, sell, or hold) for the asset under consideration. These tasks are persistent, in that they are continuously present. An agent will be dealing with many such tasks simultaneously. Gathering opinions from the area experts (fundamental analysis, technical analysis, news, and the opinions of other analysts—the published output of similar human organizations) requires registering with them and then either waiting for new opinions to be received or asking for them directly. An opinion consists of not just a buy/sell/hold recommendation but a short list of positive and negative reasons for holding that opinion, and potentially both symbolic and numeric measures of uncertainty. Information integration involves removing redundant information, resolving conflicts (or declaring them unresolvable), and forming a coherent group opinion, that can then be used for decision-making (in the light of the user’s risk tolerance, investment goals,

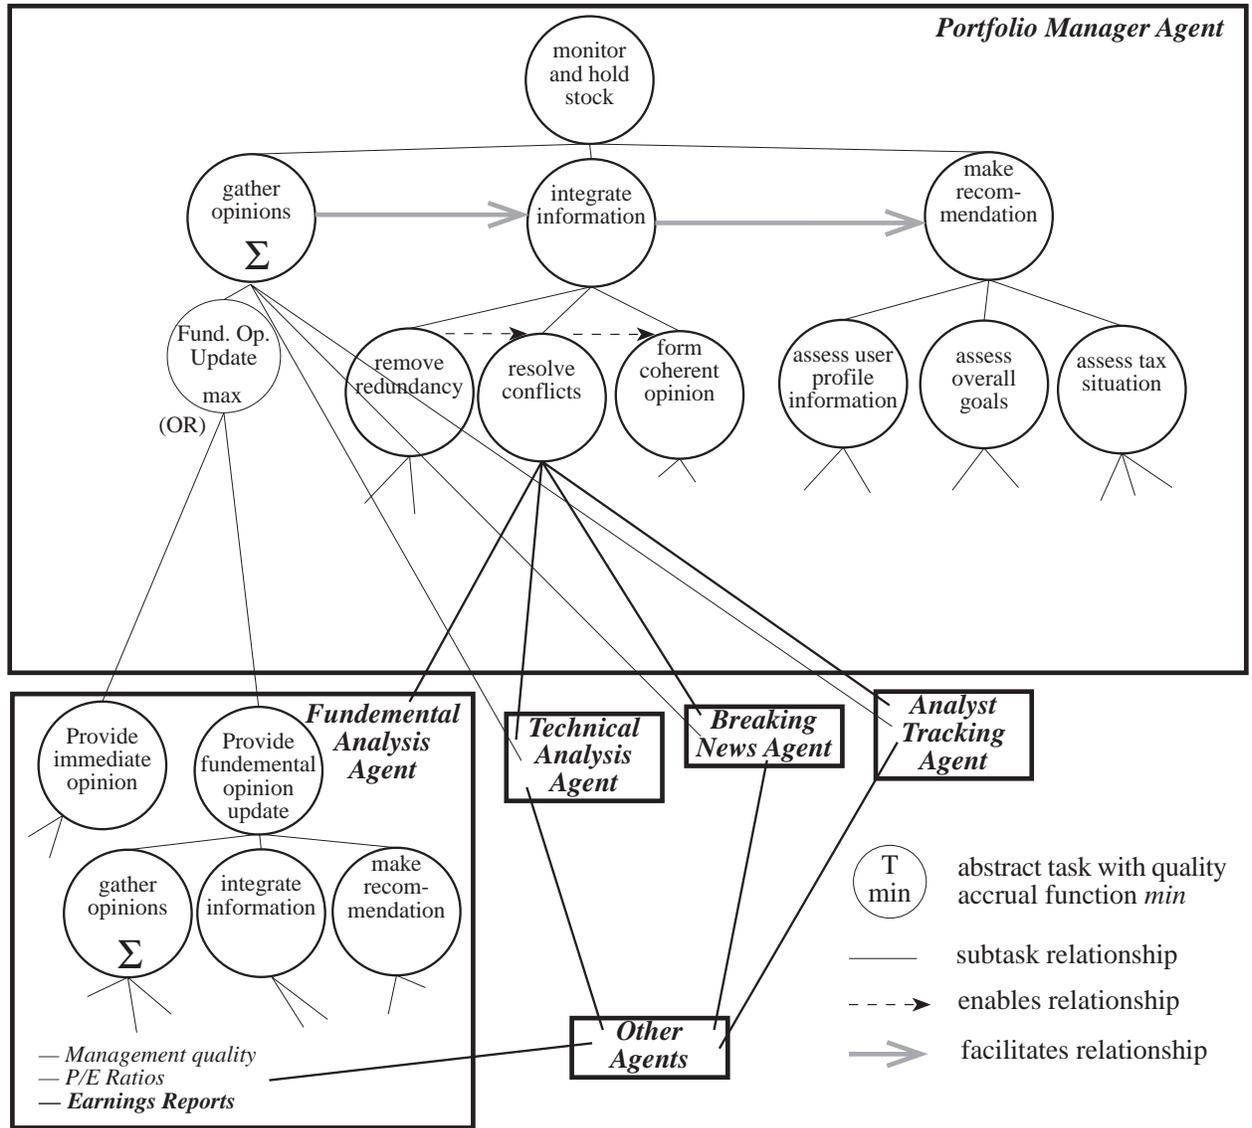


Figure 2: A task structure representing high-level portions of the monitor-stock-and-hold task.

asset allocation, tax status, and so on). The conflict resolution process may involve negotiation between the agents involved.

4.2 An Example of Coordination: Earnings Report Interpretation

One interesting subproblem in portfolio management is acquisition and interpretation of earnings reports and earnings estimates. The earnings analysis task is a complex one that includes estimating the impact of one company's earnings on other companies in a sector, the information contained in one company's earning report that actually releases information about all companies in a sector, timing in the release of earnings reports (especially for smaller companies), and the differences in actual earnings versus expectations. It is important to track revisions in earning estimates over time, as they often give important clues as to future price moves.

Figure 3 shows a relationship between the abstract plans of the earnings analysis agent and the human analyst opinion tracker agent. The earnings analysis agent initially needs to get data on a companies current and historical earnings patterns, and then it needs to keep up to date on new earnings reports as they are released. Not only does it need to track the new earnings of the company in question, but also the earnings of other companies in an industry sector. For example, a change in the portion of earnings attributed to sales is often applicable to all companies in a group, unlike changes to costs (sales minus earnings) [7].

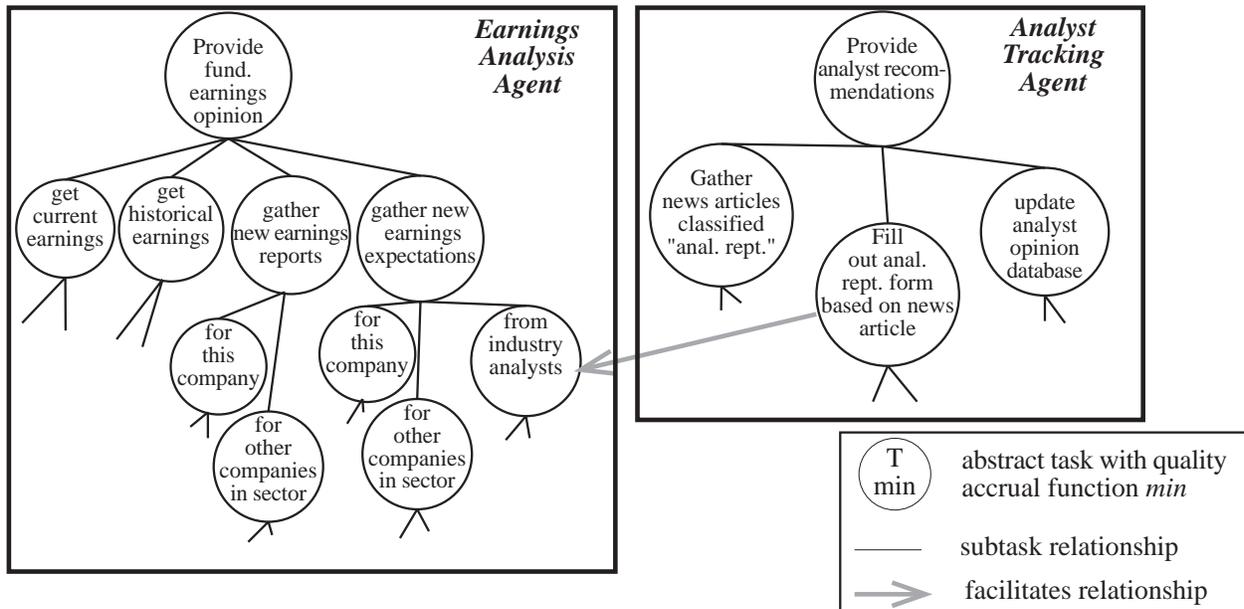


Figure 3: A task structure showing one coordination relationship between tasks in the domain of the earnings report agent and the human analyst tracking agent.

The analyst tracking agent gathers, from news and other sources, existing and updated analyst reports on a company, including revised earnings estimates (often part

of a larger report). This part of the data, if transmitted to the earnings analysis agent, can somewhat speed up (i.e., facilitate) the process of gathering earnings expectations. We have demonstrated the use of general coordination mechanisms, called the GPGP (Generalized Partial Global Planning) approach, that can easily coordinate such task structure interactions [3]. In this instance, the soft-predecessor-coordination-relationship mechanism will cause the analyst tracking agent to commit to the transmission of a completed analyst report form to the earnings analysis agent, which can then easily extract the portion dealing with the updated earnings estimate.

5 Agent Architecture

The portfolio manager and task assistant agents have an internal agent structure called DECAF (Distributed, Environment-Centered Agent Framework)—a general, reusable, core agent control architecture. The term *architecture* here refers to the internal control structure of a single agent, as opposed to the term *organization* that refers to the control and communication structure of a group of agents.

The important features of the DECAF architecture are:

- A set of clearly defined control modules (planning, coordination, scheduling, decision-making, and monitoring) that work together to control an agent. The implementations of these modules are not fixed, only their interfaces are. It is often helpful to imagine these modules as acting asynchronously (or multi-threaded) even if they are not implemented in this manner. Each module has an essentially null implementation, so agents don't have to be any bigger than they need to be. The idea is that properly developed modules can be re-used and traded.
- A core task structure representation that is shared by all of these control modules. This core structure can be annotated and expanded with all manner of details that might be “understood” only by one or a few control modules, but there is a core, shared representation.

Briefly, the main control functions consist of a *planner* that creates or extends the agent's view of the problem(s) it is trying to solve, called the *task structure*. The *coordinator* notices certain features of that structure, and may annotate it, expand it, communicate parts to other agents, or add scheduling constraints to it. The *local scheduler* takes the rough plan and creates a low-level schedule or schedules that fix the timing and ordering of actions. The *decision-maker* chooses a schedule—if there is more than one—that best meets the agent's current needs or performance goals. The *execution monitor* takes care of actually executing the next desired action (perhaps including pre-emption of the action in true real-time execution).

Previous work has focussed on the design of the coordinator [3], the local scheduler[6], and how they interact with the decision maker[5]. Details of the implementation of these components can be found in the cited papers. We intend to extend this work to include more sophisticated execution monitoring, using such techniques as the TCA (Task Control Architecture) approach [13].

6 Conclusions

We have presented the overall framework and design decisions made in our multi-agent system for the management of financial portfolios through information access, filtering and integration. Within this framework we will explore research issues of agent coordination and negotiation and case base structuring for user, task and situation modeling. In addition, there are a number of learning-related research issues we want to explore. How do we formulate the learning task in the context of multi agent interactions where procedural and control knowledge must be learned? Concept learning has been the focus of most machine learning research (e.g., [10]). Learning of control knowledge has been explored using case-based reasoning (e.g., [8, 20, 11]), and reinforcement type learning techniques (e.g., [14, 12]). This research has been conducted almost exclusively in a single agent setting. We want to explore strategies for multiple agent learning of control knowledge during agent interactions. Within each formulation of the learning task (e.g. as a case-based learning, or reinforcement learning), there are additional more specific issues to be explored. For example, for multiple agent case-based learning, new case indexing and retrieval algorithms might be necessary. In addition, the number of training cases that must be incrementally acquired through agent interactions for reliable learning is an open issue.

References

- [1] Keith S. Decker. *Environment Centered Analysis and Design of Coordination Mechanisms*. PhD thesis, University of Massachusetts, 1995.
- [2] Keith S. Decker and Victor R. Lesser. Quantitative modeling of complex computational task environments. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 217–224, Washington, July 1993.
- [3] Keith S. Decker and Victor R. Lesser. Designing a family of coordination algorithms. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 73–80, San Francisco, June 1995. AAAI Press. Longer version available as UMass CS-TR 94–14.
- [4] Oren Etzioni and Daniel Weld. A softbot-based interface to the internet. *Communications of the ACM*, 37(7), July 1994.
- [5] Alan Garvey, Keith Decker, and Victor Lesser. A negotiation-based interface between a real-time scheduler and a decision-maker. In *AAAI Workshop on Models of Conflict Management*, Seattle, 1994. Also UMASS CS TR–94–08.
- [6] Alan Garvey, Marty Humphrey, and Victor Lesser. Task interdependencies in design-to-time real-time scheduling. In *Proceedings of the Eleventh National Conference on Artificial Intelligence*, pages 580–585, Washington, July 1993.
- [7] G. Joh and C. Lee. Stock price response to accounting information in oligopoly. *Journal of Business*, 65(3):451–472, July 1992.
- [8] Subbarao Kambhampati and James A. Hendler. A validation-structure-based theory of plan modification and reuse. *Artificial Intelligence*, 55(2-3):193–258, 1992.

- [9] H. Markowitz. *Portfolio selection: efficient diversification of investments*. B. Blackwell, Cambridge, MA, second edition edition, 1991.
- [10] Ryszard Michalski and Gheorghe Tecuci. *Machine Learning: A multistrategy Approach*, volume IV. Morgan Kaufmann Publishers, 1994.
- [11] Kazuo Miyashita and Katia Sycara. Cabins: A framework of knowledge acquisition and iterative revision for schedule improvement and reactive repair. *Artificial Intelligence*, 76(1–2), 1995.
- [12] Andrew W. Moore. Prioritized sweeping: Reinforcement learning with less data and less real time. *Machine Learning*, October 1993.
- [13] R. Simmons. Structured control for autonomous robots. *IEEE Trans. on Robotics and Automation*, 10(1), February 1994.
- [14] R. S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.
- [15] K. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, Detroit, Mich, 1989.
- [16] K. Sycara. Multi-agent compromise via negotiation. In M. Huhns and L. Gasser, editors, *Distributed Artificial Intelligence*, volume Volume 2. Pittman, 1989.
- [17] Katia Sycara and Dajun Zeng. Task-based multi-agent coordination for information gathering. In Craig Knoblock and Alon Levy, editors, *Working Notes of the AAAI Spring Symposium Series on Information Gathering from Distributed, Heterogeneous Environments*, Stanford, CA, March 1995. AAAI.
- [18] Katia Sycara, Dajun Zeng, and Kazuo Miyashita. Using case-based reasoning to acquire user scheduling preferences that change over time. In *The Proceedings of the Eleventh IEEE Conference on Artificial Intelligence Applications (CAIA '95)*, Los Angeles, February 1995. IEEE.
- [19] R. Trippi and E. Turban, editors. *Investment management: decision support and expert systems*. Van Nostrand Reinhold, New York, 1990.
- [20] Manuela M. Veloso. *Learning by Analogical Reasoning in General Problem Solving*. PhD thesis, School of Computer Science, Carnegie Mellon University, 1992.