Opportunistic Enterprises in Invention

Marin Simina, Janet Kolodner, Ashwin Ram
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
{marin, jlk, ashwin}@cc.gatech.edu

Michael Gorman
TCC, SEAS
University of Virginia
Charlottesville, VA 22901
meg3c@virginia.edu

Abstract

This paper identifies goal handling processes that begin to account for the kind of processes involved in invention. We identify new goal properties and mechanisms for processing goals, as well as means of integrating opportunism, deliberation, and social interaction into goal/plan processes. We focus on *enterprise goals*, which extend traditional design goals and knowledge goals to address significant enterprises associated with an inventor. Enterprise goals represent "seed" goals of an expert, around which the whole knowledge of an expert gets reorganized and grows more or less opportunistically. Enterprise goals reflect the idiosyncrasy of thematic goals among experts. They constantly increase the sensitivity of individuals for particular events that might contribute to their satisfaction. Our exploration is based on a well-documented example: the invention of the telephone by Alexander Graham Bell. We propose mechanisms to explain: (1) how Bell's early thematic goals gave rise to the new goals to invent the multiple telegraph and the telephone, and (2) how the new goals interacted opportunistically. Finally we describe our implemented computational model, ALEC (Analogical Learning by Explaining Creatively), that accounts for the role of enterprise goals in invention.

Introduction

This paper investigates the enterprises of invention. We focus on *enterprise goals* in design, which address those enterprises of an inventor resulting in the creation of novel and interesting artifacts. Enterprise goals extend design goals (e.g., Goel et al., 1997) and knowledge goals (e.g., Ram & Hunter, 1992) with new properties, providing one way to explain how creative and innovative ideas are generated, evaluated and further pursued by expert reasoners, such as inventors.

In our view, a professional inventor pursues his enterprise ideas to satisfy his curiosity (i.e., learn more) and/or to get social recognition (e.g., to get famous and wealthy). Whenever the inventor comes across a new design idea (e.g., through experimentation or social interaction) that is interesting for him and/or for society, the preliminary preconditions for posing an enterprise goal are met. But a rational inventor must also estimate his chances of success, before allocating significant cognitive resources for pursuing a new design idea. The estimation is based on the inventor's knowledge and also on his available cognitive resources. Bell's quests for the multiple telegraph and for the telephone[1] (US v. Bell, 1908) count among enterprise goals. These goals help to expand and reorganize the whole expertise of the inventor, by providing a reason to try new approaches and learn more when classical design methods don't work. In particular, enterprise goals may generate broad knowledge goals to get domain expertise, while the knowledge gained through knowledge goals may be used to generate new enterprise goals.

Due to their high risk, and to the possibility of cross-fertilization among overlapping goals, an inventor tends to pursue several enterprise goals in parallel. Gruber (1974) used the term *network of enterprises* to describe the way scientists such as Darwin pursued a set of related enterprise goals. However, since the successful pursuit of an enterprise goal may often demand most of the recognition and processing capabilities available to the inventor, the number of enterprise goals pursued in parallel tends to remain low.

Our exploration is based on a well-documented example: the invention of the telephone by Alexander Graham Bell. We started by analyzing Bell's invention process in terms of themes[2], goals and plans (Schank & Abelson, 1977), but we realized that some of the thematic goals[3] got refined and achieved a preferential status in an unexpected way, that further guided the whole invention process. As an example, Bell became obsessed by the idea of the telephone, even if it was not in his thematic expertise domain (i.e., acoustics).

Social interaction played an important role for the invention of the multiple telegraph and telephone. Bell needed Watson's help to build, evaluate and repair electrical artifacts. Through collaboration, Bell learned specific plans to achieve his enterprise goals. Reading electricity books and performing electrical experiments were among those plans. But it was the opportunistic recognition of information interesting to his enterprise goals that guided Bell's learning processes. Unlike traditional goals, enterprise goals may remain active after finding several alternative design solutions[4].

Whenever Bell came across a better model/design in the service of an enterprise goal (suspended or active in the background), he learned/assimilated the new alternative. As an example, Bell's initial conceptual model for the telephone (i.e., the harp apparatus[5]) was based on the idea that speech must be decomposed explicitly in its harmonic tone constituents for electrical transmission at a distance. Later, during an experiment with the multiple telegraph, Bell came across a new design able to transmit all the components of speech at a distance without decomposing it explicitly into simple tones. At once, the goal of inventing the telephone was opportunistically remembered, and the new design was considered as a

[1]Alexander Graham Bell's Notebooks are available on the WWW at: *http://jefferson.village.virginia.edu/~meg3c/id/albell/homepage.html*

[2]A *theme*, according to Schank, is a generator/predictor of thematic goals.

[3]*Thematic goals* are those goals that are "around" all the time, and can be predicted for an individual.

[4]Old design solutions provide a context for adapting and evaluating new solutions.

[5]The *harp apparatus* was Bell's mental model for a device that could transmit either musical tones or speech.

more promising solution for it.

Whenever a recently considered/processed artifact was primed during the elaboration of the current enterprise goal, that artifact was considered as a potential alternative to satisfy the enterprise goal. For example, while working on the telephone microphone, Bell was also performing experiments with the ear phonautograph, a mechanical device for visualizing speech. When Bell tried to design a device able to move a piece of steel in the way that the air was moved by the action of the voice, he was reminded that the ear phonautograph provided the required behavior[6].

This idiosyncratic sensitivity for enterprise goals behaved like a kind of "knowledge lens", which helped Bell to deal with large amounts of information, by focusing him only on the relevant parts. The relevant information was constantly reorganized and learned as new cases (i.e., chunks of knowledge). Consequently, Bell achieved a goal-directed expertise in electricity, very efficient for his goals, but which sometimes ignored the traditional view of the domain.

Based on our analysis of Bell's inventions, this paper identifies a process of goal/plan handling that begins to account for the kinds of goal processing that inventors and expert researchers do. Enterprise goals extend the properties of clasical design goals. We identify new mechanisms for processing such goals, as well as means of integrating opportunism, deliberation, and social interaction into goal/plan processes.

Based on Bell's case study, we have developed a computer program, ALEC, which accounts for the role of enterprise goals and social interaction in invention and creative design. Our computer model extends the memory architecture presented in Simina and Kolodner (1995) and Simina and Kolodner (1997).

Cognitive Issues in Invention

How do interesting ideas give rise to new enterprise goals (i.e., goals to invent new artifacts)? What knowledge and processing is relevant for pursuing an enterprise goal? What is the role of social interaction in invention? These are some of the issues that a cognitive model of invention must address. Our analysis of Bell's quest for the multiple telegraph and telephone has identified ways in which existing goal processing methods (e.g., Schank & Abelson, 1977; Kulkarni & Simon 1988; Hammond et al., 1992; Ram & Hunter 1992) must be modified or augmented to handle invention.

A Critical review

We started by analyzing Bell's reasoning in terms of goals, plans and themes, according to Schank and Abelson's (1977) computational model of goal generation and refinement. Using Schank and Abelson's model, we could easily identify Bell's main *life theme* as a "teacher of the deaf", which generated the goal to invent machines that would make it easier for the deaf to hear and to learn to speak (e.g., to "visualize speech" by providing visual feedback). But we could not identify a straightforward theme-goal-plan chain to account for Bell's goal to invent the multiple telegraph or the telephone. After all, Bell was an expert in acoustics and his thematic goals, as taught by his father, had nothing to do with electricity. Nor is it clear why a teacher of the deaf would want to invent devices such as a multiple telegraph or telephone.

[6]Note that the required behavior was not reflected in the phonautograph function, namely to transform the speech in a graphical representation. Consequently, Bell could not remember the phonautograph based solely on its function. This issue is addressed in Simina and Kolodner (1995).

According to our analysis, the generation of enterprise goals is a *deliberative* process, in which the rule-based mechanisms for goal generation proposed by Schank and Abelson (1977) are only part of the story. Even if an inventor finds some design ideas interesting for himself and/or society, he will not decide instantly to pursue these ideas. Before that, an inventor must also estimate if he has the prerequisite knowledge and cognitive resources to pursue the idea successfully. Bradshaw and Lienert (1991) argue that successful estimation must be guided by a functional analysis[7] to focus and guide the further design process. But analyzing functional constraints and identifying efficient methods to further investigate those constraints is usually not a straightforward process to be completed in one reasoning session. An inventor might need to postpone the generation of an enterprise goal until he opportunistically learns methods to investigate functional constraints critical for fulfilling that goal.

Once an enterprise goal is generated, what processing mechanisms are relevant for pursuing it further? When an inventor decides to pursue an invention, he allocates significant cognitive resources to it. Even after processing, enterprise goals still remain active in background for a while, affording subtle priming effects. By strategically activating and suspending the current enterprise goal, an expert reasoner may maintain an interrelated network of active goals in the background. Similar reasoning processes were identified by Gruber (1974) in Darwin's work:

> The fact that he was all these things [the pigeon fancier, the evolutionist, and the materialist] at once meant that a unique and productive intersection of many enterprises could occur in his thinking. At the same time, the existence of this ensemble was not an accident but the deliberately cultivated fruit of Darwin's work.

Research on predictive encoding (Hammond & al., 1992) grew up from the difficulties associated with managing active goals (e.g., it is unlikely that all the goals are active due to computational demands). Basically, goals that cannot be satisfied immediately are associated at the time of encoding in memory with features of the environment in which goal achievement would likely to be possible. But in invention the structure of goals is more complex and it is difficult to enumerate, in advance, all the features of the environment in which goal satisfaction might be possible. Fortunately, we can maintain a small number of goals active at any time, without excessive computational demands. Simina and Kolodner (1995) presents a memory model which accounts for the existence of active goals and postulates their limited number. We suggest that Gruber's networks of enterprises have active-goal properties.

While the explicit purpose in invention is designing novel and useful artifacts, learning also plays an important part. An inventor often must learn new concepts (sometimes by consulting others; sometimes by experimentation and exploration) in order to move forward with designing. Ram and Hunter (1992) describe the role of knowledge goals to guide inference and learning, in the context of story understanding and problem solving. But in addition, inventors (and scientists) have a specific way of addressing such knowledge goals through deliberative experimentation and exploration (Gorman, 1997). Indeed it sometimes seems that experimentation drives invention. KEKADA (Kulkarni & Simon, 1988) illustrates this. It proposes a scientific discovery method driven

[7]The Wright brothers functionally decomposed the flight of birds and matched it with bicycles; Bell decomposed speech transmission and matched it with telegraphic equipment.

by experimentation. But in allowing itself to be driven by experimentation, KEKADA misses opportunities to achieve its goals. Namely, KEKADA ignores its background goals while it focuses on a current experimentation subgoal. A combination of deliberative and opportunistic processing of goals is needed.

The role of social interaction has, in general, been minimized in cognitive models of invention. Both IMPROVISER (Wills & Kolodner, 1993) and IDeAL (Bhatta & Goel, 1993) propose an oracle for modeling external interaction. In IDeAL, an oracle supplies the needed information when the system itself fails to solve a problem and the system assimilates the new information so as not to fail for the same reason in future. Thus IDeAL's interaction with other agents is quite limited. Our analysis of Bell's notebooks suggests that external interaction in invention takes the elaborated form of social interaction. An inventor may request information or assistance from his peers, but an answer is not guaranteed, nor is it guaranteed in a timely way. New processes are responsible for handling social interaction.

Enterprise Goals in Invention

To deal with the above invention issues, we identified a distinguished class of goals, namely *enterprise goals*. Enterprise goals are open goals, outside the operational knowledge of a reasoner, which are "interesting" for him, and for which the reasoner has some "competence" to approach them. Enterprise goals borrow from design goals (e.g., Goel et al., 1997) and knowledge goals (e.g., Ram and Hunter, 1992), but they: (1) fulfill explicit thematic goals of a reasoner, (2) have a long-term significance and activity assessed explicitly (enterprise assessment), (3) rely explicitly on opportunistic cross-fertilization processing (even partial results may help satisfy overall thematic goals indirectly), and (4) are actively seeking solutions by planning and acting (e.g., experimentation), filling implicit knowledge gaps in the process. Enterprise goals are usually long-term goals, are not easily satisfied (they might look for several alternative solutions to facilitate reasoning and evaluation), and they require both theoretical qualitative reasoning and experimentation for their achievement.

How do interesting ideas give rise further to new enterprise goals? In our view, an idea is interesting for an inventor when: (1) it is instrumental to satisfying some of his higher thematic goals, and (2) he has some idea about how to investigate it. Some of the higher thematic goals are personal (e.g., scientific curiosity), while others have to do with what society will value. If an invention idea is judged to be interesting for the inventor and/or the society, the inventor must also deliberate if he has a minimal knowledge (e.g., functional decomposition to quickly assess what is known) and the cognitive resources for pursuing it further. Enterprise goal generation is a *deliberative* process in invention.

Let's illustrate the above ideas with the episode when Bell decided to approach the invention of the multiple telegraph. In October 1872, Bell read a newspaper article[8] describing the impact of a new telegraph system able to transmit simultaneously two telegraphic messages over the same wire (Stearn's duplex system). The article also suggested that fame and fortune awaited the inventor of a telegraphic system able to transmit more than two messages simultaneously. The multiple telegraph idea was definitely interesting for society, and this could be an incentive for Bell to check the relevance of the multiple telegraph idea for his own research.

[8] In *Boston Transcript*, the newspaper in which Bell advertised his speech lessons.

Bell functionally assessed the problem as follows: (1) multiplex multiple messages and send them over telegraphic wire, and (2) demultiplex the messages at destination. Bell remembered that he had some experience with telegraphic equipment, while he was trying to understand Helmholtz's Apparatus[9]. According to Bell's understanding, Helmholtz's Apparatus was able to unscramble (demultiplex) multiple tones sent over a single wire, by using tuned receivers. Since both devices performed demultiplexing, Bell could use his understanding of Helmholtz's Apparatus for the multiple telegraph idea. Bell estimated that he also had the prerequisite knowledge for inventing the multiple telegraph. Helmholtz's Apparatus provided an easy solution for half of the problem (demultiplexing), while the multiplexing part looked even easier: just add several tone generators in series with the one existing already in Helmholtz's Apparatus. Given these premises, Bell decided to pursue the invention of the multiple telegraph (as an enterprise goal), for which he already had a theoretical model and a partial implementation.

What knowledge and processing are relevant for pursuing enterprise goals? Enterprise goals are goals that identify interesting artifact ideas (functions and behaviors), to be pursued further as design goals. Consequently, enterprise goals may contain a behavioral device specification, possibly incomplete and inconsistent. The synthesis of a structural solution may be facilitated by evolving in parallel several alternative solutions to the design specification (Wills & Kolodner, 1993; Gruber, 1974). As a side effect of pursuing enterprise goals, the expertise of the inventor increases. Namely, if enterprise goals cannot be pursued due to the lack of expertise, they may spawn broader knowledge goals to get that expertise.

In our view, two complementary and interacting processes are responsible for incrementally evolving an invention (Kolodner & Wills, 1993): (1) Enterprise Processing, which proposes new design solutions to the current design specification, and (2) Evaluation & Repair, which critiques the current design solution and may update the design specification. Both processes rely heavely on the previous experience of the inventor. Enterprise Processing relies on a library of known artifacts and techniques for adapting them to fit the current design specification, while Evaluation & Repair relies on simulation, experimentation and knowledge to interpret the experimental results. The critique provided by Evaluation & Repair may also suggest a divide and conquer strategy for pursuing an enterprise goal by decomposing it into subgoals. Gorman (1992) reports that Bell used such a strategy to independently elaborate critical functional subparts of the telephone.

But what processing is responsible for generating such subgoals? Our exploration of the invention of the multiple telegraph and of the telephone suggests the following approach. The inventor chooses a design alternative responsible for satisfying the main, or most difficult to satisfy, constraints. Then he generates subproblems to satisfy the other secondary constraints, in the framework of the main design alternative. For the invention of the telephone, the main constraint was to transmit an "undulatory current" (i.e., both pitch and amplitude) over a telegraphic wire. Once Bell identified such a design, he generated subproblems to improve its secondary characteristics. Those subproblems resulted in the design of the telephone's microphone and receiver.

Let's also address some of the processing differences between enterprise goals and thematic goals. While themes

[9] Helmholtz's Apparatus was an electrical device for producing artificial vowels, relevant to Bell's acoustical research.

may characterize the goals common to most inventors, they don't say anything about how inventors *idiosyncratically* refine these themes and why different individuals working in the same domain may chose different paths to achieve their (similar) thematic goals. Also, the traditional view of thematic goals does not explain why and how inventors allocate more computational resources to some thematic goals (keeping them more active), but not to others. Enterprise goals are intended to explain the idiosyncrasies of individual experience among experts, providing a "knowledge lens" for efficient interaction with the events noticed in the world. Our analysis of Bell's diaries and related literature suggest the following mechanism. Once an expert evaluates the potential of a goal to fulfill at least one of his major thematic goals, he activates and focuses mostly on that (enterprise) goal. The expert ignores, for a while, other ways to satisfy his thematic goal, since the achievement of an enterprise goal will implicitly satisfy the associated thematic goal.

What is the role of social and environmental interaction in invention? In our view, a cognitive model of invention must account for social and environmental interaction. Otherwise, the model fails to give a plausible account of the huge computational resources involved in acquiring the "right" knowledge for making the invention possible. A real expert takes advantage of the knowledge and processing available elsewhere through strategic social and environmental interaction.

An expert inventor is part of a society of experts (agents), who may communicate among themselves. Each expert has similar cognitive abilities, but different knowledge. Consequently, particular enterprise goals may be addressed by some experts, but not by others. Each expert can ask (i.e., make queries) to: (1) retrieve knowledge available elsewhere, and (2) request external evaluation of a proposed design. Once an expert learns to design a new artifact, all of the other experts in the society have access to that design through the publications and social interactions of that expert.

By reading a newspaper article (i.e., environmental interaction), Bell learned about the importance of inventing the multiple telegraph. The Enterprise Assessment process found the problem interesting and feasible, so it generated the goal to invent the multiple telegraph. What about the role of social interaction? If the Enterprise Processing process cannot propose a design solution to the current specification, it may ask other experts for a solution. In particular, this was the reason why Bell hired Watson for developing electrical artifacts. Also, if a design solution proposed by the Enterprise Processing process cannot be evaluated/repaired by the Evaluation & Repair process (due to his lack of previous experience), it may ask other experts to perform the evaluation. For example, when Bell had doubts about the physical principles involved in telephony, he requested help from distinguished scientists, such as Joseph Henry.

Computational Model

Architecture

Our computational architecture for invention is implemented in a computer program called ALEC (see Figure 1). We assume ALEC is an agent, representing an expert inventor, that can collaborate with other agents (same architecture, but different knowledge).

Since our view of the invention process relies heavily on event detection and processing, we opted for an event-driven architecture. A WORKING MEMORY (WM) keeps track of all the state information which may generate internal or external events. Since the processes (represented as gray rectangles in Figure 1) operating on the WM rely heavily on previous experience, we also need a LONG-TERM MEMORY (LTM) to account for the role of experience in invention. Our model builds on the blackboard model of WM (Hayes-Roth 1985) and on the CBR memory model of LTM (Kolodner, 1993). Basically, changes in the WM generate events used to activate or generate new goals in the AGENDA. A Strategic Control process selects the next CURRENT GOAL from the AGENDA, based on a CONTROL PLAN.

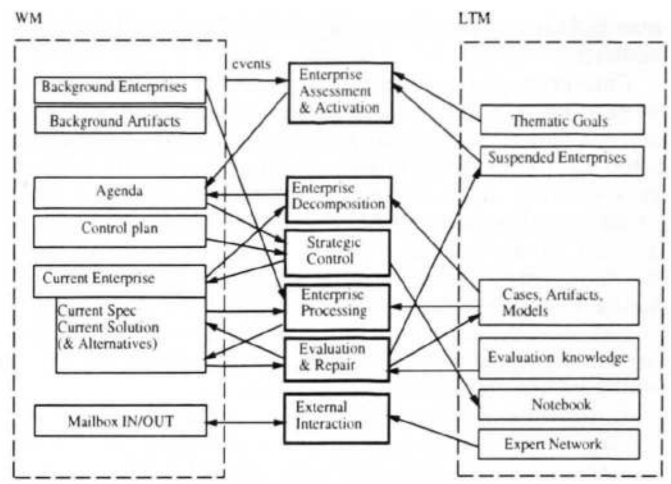

Figure 1: ALEC: a Framework for Invention

Let's elaborate the LTM of our architecture. In an event-driven architecture, the monitored events may remember SUSPENDED ENTERPRISES and THEMATIC GOALS. The SUSPENDED ENTERPRISES are indexed in LTM in terms of their *preconditions*, represented as monitored events. To facilitate the generation of design solutions, given their intended function, the LTM contains a collection of CASES, ARTIFACTS, MODELS, indexed by their function. To facilitate internal evaluation of the CURRENT SOLUTION proposed by the Enterprise Processing, the LTM contains EVALUATION KNOWLEDGE (e.g., Forbus, 1984) used by the Evaluation & Repair. To facilitate a further analysis of the invention process, LTM contains a NOTEBOOK, which provides a derivational record of the invention.

Once an expert learns about other relevant experts (and their expertise domains), its External Interaction process must be able to use this knowledge, if needed, to send request messages. Consequently, LTM should be also a repository of knowledge about EXPERT NETWORK, indexed by their expertise domains.

The WM represents the "activated" part of LTM. To simulate *priming effects*, WM keeps track of the recently processed design goals, BACKGROUND ENTERPRISES, and of the recently processed artifacts, BACKGROUND ARTIFACTS. A decay process (not represented in the Figure 1) limits the number of activated items in WM. Based on the events noticed in WM, the Enterprise Assessment & Activation may update the AGENDA by: (1) activating suspended enterprise goals, or (2) generating enterprise goals, instrumental for satisfying thematic goals. The Strategic Control selects a CURRENT ENTERPRISE from the AGENDA. The CURRENT GOAL may contain a CURRENT SPEC of a design (possible incomplete and inconsistent), a CURRENT SOLUTION, and a set of ALTERNATIVES (used by the Enterprise Processing to propose new designs). The MAILBOX IN/OUT is used for communication with other experts, via External Interaction. If a design goal cannot be achieved using the knowledge available locally, a request is sent to relevant experts. Since an answer is not guaranteed and it may come asynchronously, the system must be able to predictively encode the goal in memory and to approach other enterprise goals from the AGENDA in the meantime.

Algorithms

Now we can describe the main algorithms for creating and manipulating enterprise goals, essential to our theory of invention.

Enterprise goal generation is part of the Enterprise Assessment & Activation process. Whenever ALEC learns about a new design idea (by social interaction or experimentation), it performs the following steps: (1) identify if the idea is *interesting* (does the implementation of the idea result in conditions that match those of important unsatisfied thematic goals?), (2) estimate its expertise to implement the idea as an artifact (given the idea's rough design spec, can the Artifactor Generator generate an artifact, judged/simulated as promising by the Evaluation & Repair?) (3) generate an enterprise goal (instrumental for satisfying specific thematic goals), if the results of steps (1) and (2) are positive.

Enterprise Processing attempts to derive new designs by opportunistically adapting and merging pieces of known designs. This involves WM priming and LTM retrieval of devices based on a given design specification according to the following enterprise processing algorithm: (1) if the current enterprise specification matches any pieces (behavioral segments) of the design alternatives associated with the Background Enterprises, adapt (by analogy) those pieces and add them to alternative solutions of the current enterprise (2) otherwise, if LTM retrieval is successful, extend the alternative solutions of the current enterprise with the retrieved devices, and check also if the behavior of those devices is opportunistically relevant for the background enterprises, (3) otherwise, send a message to other experts for artifacts fulfilling the enterprise specification, and suspend the enterprise in memory, indexed by its design specification. The Enterprise Processing algorithm is responsible for the synergy among (related) enterprises evolved in parallel.

In Figure 2 we have represented two enterprises, Enterprise 1 and Enterprise 2, which evolve in parallel. A functional specification of an enterprise accounts for an initial *given* state, which is transformed by the device function in a final *makes* state (Goel et al. 1997). Each enterprise specification contains a *makes* state (e.g., M1 for Enterprise 1), a *given* state (e.g., G1 for Enterprise 1), and possibly, constraints restricting the (behavioral) paths between the given and makes states (e.g., C1 for Enterprise 2). The design synthesis task is responsible for incrementally evolving a path between the *given* and *makes* state, relying on analogy. In the CBR framework, the adaptation step is responsible for such behavior.

If the Enterprise Processing is unable to retrieve artifacts that satisfy most of the design spec constraints, it attempts Enterprise Decomposition of the initial enterprise into subgoals to facilitate the further synthesis of the desired artifact. The algorithm for functionally decomposing an enterprise goal into subgoals, given the design specification of the artifact is: (1) find a design alternative that satisfies the main constraints[10], (2) identify which components of this design alternative are responsible for satisfying the secondary constraints, (3) generate subproblems for designing components, to satisfy (better) the secondary constraints.

[10]The main constraints are most difficult to satisfy (i.e., a very limited number of retrieved artifact cases may satisfy them).

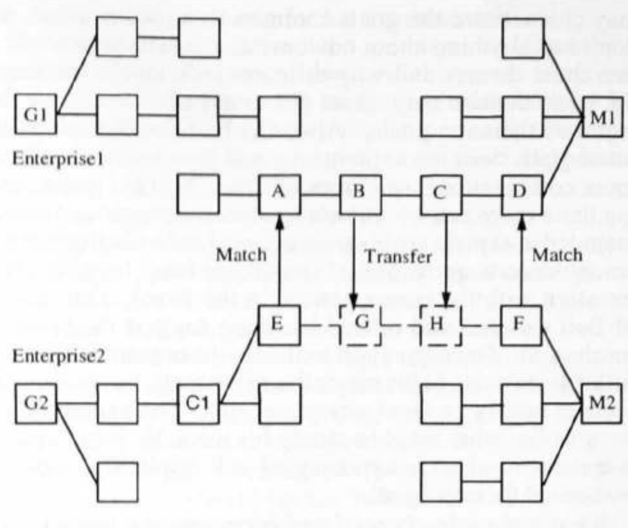

Figure 2: Enterprise Synergy

But Evaluation is essential to validate the design solutions proposed by the Enterprise Processing. Here are the steps: (1) perform simulation of the artifact and critique it, if ALEC has enough domain knowledge, (2) otherwise, implement the artifact, perform experiments and interpret the results, if this is possible, (3) otherwise, send message to relevant experts to remotely perform the evaluation.

Step (3) of the Enterprise Processing retrieval and Evaluation & Repair algorithms may result in the deliberative generation of broad knowledge goals to provide missing expertise. The details are beyond the scope of this paper.

Example: the Multiple Telegraph Idea

After reading a newspaper article suggesting that fame and fortune awaited the inventor of a multiple telegraph, Bell realized that he could design a multiple telegraph based on his knowledge of Helmholtz's Apparatus and acoustics. Moreover, the goal to invent the multiple telegraph became Bell's enterprise goal for the next several years. We present a possible analysis of Bell's reasoning, in terms of our computational model (ALEC) presented in Figure 1.

Reading the newspaper article about the multiple telegraph is equivalent in our system to receiving a message (in the MAILBOX), containing a design goal (i.e., Multiple Telegraph), characterized by a sketchy DESIGN SPECIFICATION and a postcondition describing that the inventor will become *famous & wealthy*. This event is analyzed by Enterprise Assessment, which identifies that *famous & wealthy* is a postcondition associated with some of Bell's thematic goals (i.e. designing the multiple telegraph might provide a *novel* alternative plan to satisfy thematic goals) and is also a precondition to other thematic goals (e.g., *get married*). Consequently, the Multiple Telegraph is *interesting*, and if ALEC could relate it to its own expertise, the Multiple Telegraph would become a good candidate for an enterprise goal. Based on the specification of the *Multiple Telegraph* (i.e., multiplex telegraphic messages over a single wire and demultiplex them at destination), Enterprise Processing searches the memory of CASES, ARTIFACTS, MODELS and retrieves Helmholtz's Apparatus, which was able to demultiplex a collection of harmonic tones. After a quick case-based *adaptation*, involving using several generators instead of only one, Enterprise Processing has a theoretical design solution. Evaluation theoretically simulates that indeed the new adapted design satisfies the CURRENT SPEC, and it reports the result to Enterprise Assessment, which promotes the Multiple Telegraph as an enterprise goal, by associating with it an initial (theoretical) solution. The new enterprise goal is put in the Agenda and will be further selected for processing.

Discussion

ALEC, our model for invention, addresses some of the control problems associated with the invention processes. In our opinion, a realistic computational model for invention should emphasize both the deliberative and opportunistic components of the control in invention. We intend for enterprise goals to provide appropriate mediation between the deliberative and opportunistic components of the reasoning process in invention. Namely, generating an enterprise goal is a deliberative process, intended to further guide the opportunistic evolution of the invention.

ALEC integrates and extends some of the previous models of invention and scientific discovery, such as Bradshaw and Lienert (1991), Kulkarni and Simon (1988), and Bhatta and Goel (1993). In particular, ALEC's enterprise goal assessment attempts a *functional decomposition* of an enterprise goal in subgoals (suggested also by Bradshaw & Lienert, 1991) to assess if at least some of these functional subgoals can be implemented easily, providing a basis for opportunistic design synthesis. After a design solution is opportunistically synthesized, it is theoretically evaluated (Bhatta & Goel, 1993), and further evaluation by experimentation may identify gaps in the theoretical model of the proposed device. The Repair algorithm is an opportunistic version of KEKADA (Kulkarni & Simon 1988). IDEAL (Bhatta & Goel, 1993), another model of invention, ignores the opportunistic aspect of control, and consequently has to rely on an oracle to provide "the right information at the right time". Our model recognizes the limited computational power available to an inventor, but provides something more realistic than an oracle, namely the ability to suspend and come back to goals based on experimentation and social interaction. Bell's case study shows that an inventor that acts in society must necessarily be opportunistic within the framework supplied by her or his thematic goals and resulting network of enterprises. Therefore, to use Shrager and Langley's (1990) term, our computational model takes a large step towards dealing with the problem of embeddedness.

ALEC's architecture grew up from the difficulties encountered with IMPROVISER's opportunistic control (Wills & Kolodner, 1994). Simina and Kolodner (1995) postulated the existence of "active goals", as a mechanism to explain priming effects, but it did not characterize which goals are more likely to be active. ALEC goes further, identifying enterprise goals as the most likely active goals in a cognitive model for invention.

We can't say yet that ALEC provides a complete or correct model of goal generation and processing in invention. We must still test it on a wide variety of examples (the invention of the telephone offers plenty, beside those mentioned briefly here), and we must investigate more deeply: (1) the synergy among active goals, (2) the role of context in goal generation and processing, (3) the relation between enterprise goals and traditional goals, and (4) ALEC's ability to simulate inventors with different backgrounds.

Acknowledgements

This research was funded in part by ONR Grant No. N00014-92-J-1234 and in part by NSF Grant No. DIR-8722002 and DIR-90-12311. We thank Ashok Goel, Hari Narayanan, Nancy Nersessian, Bill Murdock and Todd Griffith for their comments on this research.

References

Bhatta S. & Goel A. (1993). Learning Generic Mechanisms from Experiences for Analogical Reasoning. *Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society* (pp. 237–242). Hillsdale, NJ: Lawrence Erlbaum Associates.

Bradshaw, G. & Lienert, M. (1991). The invention of the airplane. In *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp.605–610). Hillsdale, NJ: Lawrence Erlbaum Associates.

Forbus, K. (1984). Qualitative Process Theory. *Artificial Intelligence* 24, (pp. 84–168).

Gruber, H. (1974). *Darwin on Man: A Psychological Study of Scientific Creativity* (2nd ed.). University of Chicago Press.

Goel, A.K., Bhatta, S., & Stroulia, E. (1997). Kritik: An early case-based design system. In Maher, M.L. & Pu, P., (Eds.), *Issues and Applications of Case-based Reasoning to Design.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Gorman, M. E. (1997). Mind in the World: Cognition and Practice in the Invention of the Telephone. *Social Studies of Science* 27:4.

Gorman, M. (1992). *Simulating Science: Heuristics, Mental Models and Technoscientific Thinking.* Bloomington: Indiana University Press.

Hammond, K., Converse, T., Marks, M. & Seifert, C. (1993). Opportunism and Learning. *Machine Learning*, 10 (pp. 279–309).

Hayes-Roth, B.(1985). A Blackboard Architecture for Control. *Artificial Intelligence* 26 (pp. 251–321).

Kolodner, J. (1993). *Case-Based Reasoning.* Morgan Kaufmann.

Kolodner, J. & Wills, L. (1993). Case-based Creative Design. *AAAI Spring Symposium on AI and Creativity.* Stanford, CA. March 1993.

Kulkarni, D. & Simon, H. (1988). The processes of scientific discovery: The strategy of experimentation. *Cognitive Science*, 12 (pp.139–175).

Ram, A. & Hunter, L. (1992). The Use of Explicit Goals for Knowledge to Guide Inference and Learning. *Journal of Applied Intelligence*, 2(1) (pp.47–73).

Schank, R., & Abelson, R. (1977). *Scripts, Plans, Goals and Understanding.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Shrager, J., & Langley, P. (Eds.) (1990). *Computational Models of Scientific Discovery and Theory Formation.* Morgan Kaufmann.

Simina, M. & Kolodner, J.L. (1997). Creative design: reasoning and understanding. In Leake, D.B. & Plaza E. (Eds.), *Case-based reasoning research and development, ICCBR-97.* Springer-Verlag.

Simina, M., & Kolodner, J.L. (1995). Opportunistic Reasoning: A Design Perspective. *Proceedings of the Seventeenth Conference of the Cognitive Science Society* (pp. 78–83). Hillsdale, NJ: Lawrence Erlbaum Associates.

US v. Bell (1908). *The Deposition of Alexander Graham Bell,* Boston: American Bell Telephone Company.