Service Agents for Customer Self-Service[1]

Position Paper

Dhiraj K. Pathak

Applied Engineering

Multivendor Customer Services

Digital Equipment Corporation

September 22, 1995

Contact Information:
Dr. Dhiraj K. Pathak
Digital Equipment Corporation
Digital Drive, MK01-2B9
Merrimack, NH 03054-9501

Email: pathak@yantra.mko.dec.com
Fax: (603) 884-1036

[1]This paper describes progress in one phase of in-house research in service delivery technology. The opinions expressed are those of the author only and do not necessarily reflect those of Digital Equipment Corporation.

The computer service delivery business is concerned with solving end-user problems with multi-vendor computer hardware systems and software applications. This business is undergoing a very significant change. In the past, large corporate accounts were the main-stay of the service business. The current reality is that the computer is fast becoming a universal tool, finding its way into our homes, shops, schools, and every other place imaginable. The locus of dominant computer use is shifting from a large corporate entity to the individual. Individuals can ill-afford to enter into expensive service delivery contracts. The very large number of individual computer users who need inexpensive service for their computer hardware and software systems are contributing to the commoditization of the service delivery business.

To succeed in the commodity service delivery business, characterized by huge volumes and very small profit margins, service delivery vendors cannot rely solely on extensive teams of service representatives. The cost of manual service delivery is simply unsustainable in the commodity space. Therefore, there is a pressing need for building automated systems for *customer self-service.* The underlying assumption is that it is both feasible to build effective automated service delivery systems and that the cost of creating and maintaining such systems is less compared to the costs of the corresponding human infrastructure that these systems are meant to replace. The challenge of the paradigm of customer self-service is underscored by the difficulty of engineering automated systems that satisfy these assumptions.

We have undertaken to meet this challenge by exploring the use of intelligent *service agents* as the basis for automated problem resolution systems. Service agents can be of two types, those that enable problem resolution by information retrieval and those that enable problem resolution by automated diagnosis. Problem resolution by information retrieval means that the user's problem is solved by pointing the user to a document in an information repository that contains a solution to the user's problem. Problem resolution by automated diagnosis means that the agent attempts to identify the user's problem and then offer a solution, if one is known, for the identified problem.

There are essentially two kinds of service problems. There is the break-fix class of problems where some physical component needs be replaced. Such problems perforce require dispatch of a service technician to customer site to carry out the needed repair. The second class of problems are informational in the sense that they can be solved by providing the customer with the appropriate information and no physical transaction is required. Two factors are making this class of problems the more significant of the two. One reason is that hardware component reliability continues to improve. The other reason stems from the fact that most commodity space users cannot invest time and money in formal training programs. Therefore, many problems faced by these users are the consequence of not knowing the appropriate information. Our service agents initiative is aimed at automated customer self-service for the informational service delivery problems.

Service agents can be made accessible to customers in a variety of ways including a bulletin board system (BBS), a voice response system (VRU), CD-ROM distribution, and on the Internet. A VRU interaction is clearly limited in scope by the limited bandwidth available for interaction with users. At present, bandwidth is less of an issue with the CD-ROM distribution. In contrast, bandwidth can become a limiting factor for both the Internet and BBSs. However, bandwidth is likely to continue to increase in the case of the Internet. For CD-ROM distribution, there is an associated logistics issue to consider.

The user interface to the service agents is critical to their acceptance by the customers. We are restricting attention to two simple interaction modalities. In the first case, the user is allowed to enter a problem description in free-format text. Seldom will the interaction be complete in a single open-loop exchange with a service agent. Therefore, the agent response to the initial user input is considered a starting point for a limited local search in an appropriate problem space. For example, if the agent response is a potentially relevant document, then the user is allowed the opportunity to navigate the space of documents that the agent considers related to the retrieved document. The second interaction modality is a dialog between the user and the service agent controlled by the latter. The agent queries the user about various properties that can describe potential problems in a given domain of discourse. Some queries can suggest user actions and reporting back of the observed outcome of these actions. Therefore, the agent-guided dialog is both a mechanism for allowing the user to specify problem information to the agent as well as a mechanism for the agent to suggest appropriate diagnostic actions to the user.

Two other properties of service agents are considered important in their success in meeting the requirement for automated customer self-service. The first property is the use of multimedia explanations in the solutions returned by the service agents. Many of the informational problems can be most easily dealt with through the use of pictures. For example, consider the problem of adding a new SCSI device to an existing computer system. What would otherwise take a lot of textual explanation can be easily disposed of by using an illustrative figure detailing the cable connections that must be made. At present, we are focussed primarily on the integration of text and pictures to generate effective explanations for the service agents to deliver. The second important property of service agents is their ability to adapt to the needs of the particular customers they are interacting with. Two ways to achieve interaction adaptation are to provide the agents with user profiles and to process the actual interaction a user is having with the agent in order to modify the evolving dialog with the user.

The design of a service agent for problem resolution by information retrieval begins with the creation of a useful corpus of documents. In some cases, these documents may exist as part of another information resource. In other cases, the documents may have to be created anew. The document set has to correspond to a useful subset of actual problems encountered by users. Therefore, a pre-analysis of the most frequently occurring problems in a given domain is required. The document set is then designed to cater to these problems. The next step is to design an appropriate retrieval function for the agent to use. The retrieval function is critically dependent upon the nature of the corpus and the expected user queries. To support a feedback mechanism in the user interaction with the service agent, an algorithm to create relationships among the documents in the corpus is required. These relationships help support the local search around the initially retrieved item from the corpus. Finally, a mechanism is required to modify the user dialog as a function of the evolving user interaction with the service agent.

The design of a service agent for problem resolution by automated diagnosis begins with the determination of a set of problems that the service agent can diagnose. This set is determined by monitoring the problems users encounter in a given domain in a recent time window. The subset of most frequently occurring problems is selected for the service agent to automatically diagnose. The number of problems in this set is also tempered by the amount of effort estimated to be required in building the service agent for these problems. After problem selection, a problem representation is required. Then, the diagnostic problem space is determined. This is the space that the service agent searches in order to diagnose the user's problem. In a way, this is also the space that the user is made to traverse since the user dialog is controlled by the service agent. A critical component is the control knowledge used to guide the search in the diagnostic problem space of the agent. Then, a theory is required for recognizing that the information obtained from the user is sufficient to identify a known problem as the user's problem. The agent must also have a mechanism for synthesizing multimedia explanations for the identified problem. The need for synthesis arises because the parameters of the problem can vary from one instance to the next and therefore the corresponding solution structure also needs modification. Finally, a mechanism for modifying the user dialog with the agent as a function of the evolving user interaction is implemented to make the agent responsive to specific users.

The development of the service agents is guided by an empirical evaluation of actual user interaction with them. The performance goal for service agents is set up as the automated resolution of a certain fraction of the user problems. Therefore, it is acknowledged that a significant fraction of the user problems will continue to be directed to service representatives. What the service agents are meant to do is to act as a buffer for a certain number of the user problems so that the number of service personnel required to provide service delivery remains within acceptable limits. The success of the service agents deployment is determined by the evaluation of the cost of creating and maintaining them measured against the savings realized by the automatic resolution of a certain fraction of the total user problems.