Knowledge-base Systems for Structures in Service

Systèmes de traitement des bases de connaissance pour structures existantes

Systeme für Wissensgrundlagen in Verbindung mit bestehenden Konstruktionen

Ian F.C. SMITH
Dr. Eng.
Swiss Fed. Inst. of Technol.
Lausanne, Switzerland

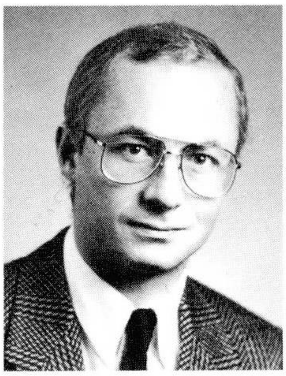

Dr. Smith received engineering degrees from Cambridge University, U.K. and the University of Waterloo, Canada. Over the past fifteen years, he has worked in research, design and construction in several countries. Presently, he is performing research into structural-engineering applications of knowledge-processing technology as well as studying various topics of fatigue and fracture in steel construction.

SUMMARY
Engineering activities related to structures in service are particularly suited to applications of knowledge-processing technology since the number of problems is increasing while appropriate knowledge remains poorly distributed. This paper examines the potential for such applications and presents a small system in order to demonstrate the advantages of rapid development of prototypes. Also, advances in artificial intelligence and improvements in hardware create a situation where hybrid reasoning techniques are feasible for many problems associated with existing structures.

RESUME
Les techniques de traitement de la connaissance s'appliquent particulièrement bien au domaine des structures existantes, car l'augmentation du nombre de problèmes n'est actuellement pas suivie d'une amélioration correspondante de la transmission des connaissances. Cet article traite les possibilités d'application de telles techniques et présente un petit systéme pour démontrer les avantages de développer rapidement des prototypes. De plus, on suggère que les progrès réalisés en intelligence artificielle et les améliorations du matériel informatique rendent possible l'application des techniques de raisonnement hybride à la résolution des problèmes liés aux structures existantes.

ZUSAMMENFASSUNG
Probleme in Verbindung mit bestehenden Konstruktionen eignen sich speziell für die Anwendung von Techniken der Wissensverarbeitung, besonders weil das entsprechende Wissen wenig bekannt ist. Der vorliegende Artikel behandelt mögliche Anwendungen und stellt ein kleines System vor, um die Vorteile einer raschen Entwicklung von Prototypen zu zeigen. Fortschritte auf dem Gebiet der künstlichen Intelligenz und der Hardware lassen hybride Beurteilungstechniken für Probleme in Verbindung mit bestehenden Konstruktionen einsetzbar werden.

1. INTRODUCTION

The number of structures exceeding their so-called design lives is increasing exponentially each year. This trend corresponds to a construction boom which began over one hundred years ago Frequently, design lives for these structures were not defined scientifically and even today, economic and political factors dominate such considerations for new structures. Historically, most design constraints have been conservative due to a lack of knowledge of material behaviour and difficulties associated with quality assurance. Therefore, most of these structures are able to remain in service well beyond their designated life.

As a result, engineers are devoting a greater proportion of their time towards tasks involving evaluation, monitoring, maintenance and modification of structures in service. Identification of the best ways to perform these tasks usually requires scientific knowledge in diverse domains, such as corrosion and fatigue, and much practical experience; traditional engineering education and standard design-office methods are rarely sufficient. Consequently, the incidence of problems related to structures in service is growing while the number of engineers skilled in identifying good solutions is limited. Therefore, such tasks are particularly suited to applications of artificial intelligence research or more specifically, development and manipulation of knowledge bases.

In spite of this opportunity, few applications of knowledge-processing technology have been implemented for structures in service. Civil engineering is a fragmented and necessarily conservative field where technological developments are not embraced blindly. Applications are hindered by factors such as uncertainty related to important parameters and difficult economic, political and social considerations.

For over ten years, other fields have experimented with technology which would be appropriate for structures in service. For example, an operational system for diagnosis in the car manufacturing industry is capable of saving tens of millions of dollars annually [1]. Specialized methods and heuristics can now be distributed widely and practical applications of machine-learning methods are becoming feasible. Also, recent developments in artificial intelligence have provided tools which may be useful to problems encountered by structures in service.

This paper identifies areas where knowledge-base systems could be applied to structures in service and reviews those systems which have already been developed. Problems typically associated with structures in service are discussed. A prototype system called CRACK CONTROL is presented. Finally, the application of more complex systems for large problems is examined considering the requirements of structures in service and recent developments of methods in artificial intelligence.

2. MANAGEMENT OF STRUCTURES IN SERVICE

Activities associated with the management of structures in service are shown in Figure 1. Over their lifetimes, structures are subjected to monitoring, evaluation, maintenance and perhaps, modification. All of these activities could benefit from efforts to organize relevant knowledge.

Monitoring in the form of regular inspections is probably the most common activity. Inspection personnel need to know what to look for, where to inspect, how to examine and how to report their findings. Knowledge-base systems can assist in such decision making, e.g. [2]. Some findings may justify an increased inspection effort. The accuracy of inspection techniques should be compatible with the sophistication of methods used to evaluate findings. Structures may also be controlled using gauges and sensors for purposes such as damping and energy conservation. Such control is relatively new and therefore, software systems which enhance performance are scarce. The number of applications will grow with increasing long-term reliability of monitoring equipment.

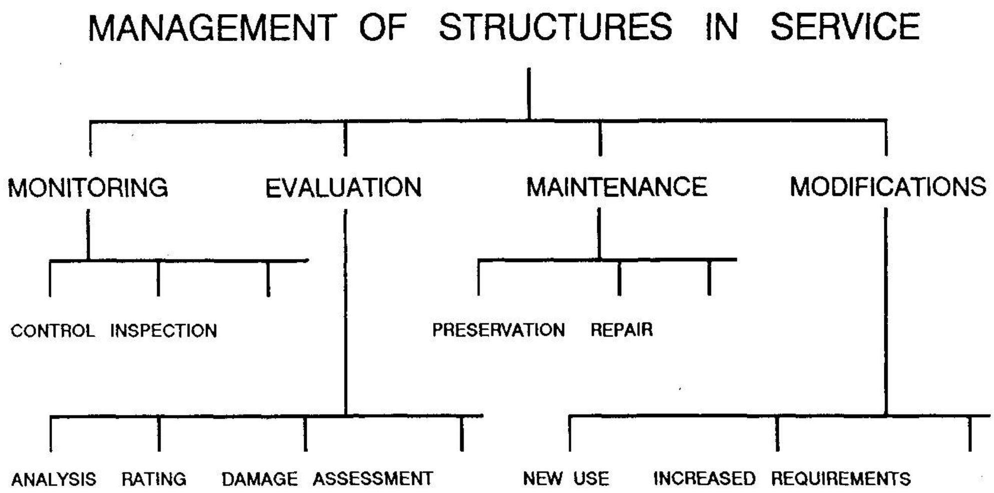

FIGURE 1 *Activities associated with the management of structures in service*

Evaluation of existing structures is the activity where the greatest effort in knowledge processing has been concentrated [3][4][5]. This concentration is understandable since often, a successful management strategy can be traced to a high quality evaluation. Applications include failure analysis, risk analysis, bridge rating and damage assessment. Analysis of structures for risk and damage due to earthquakes is thus far, the most studied domain. Other criteria for evaluating structures, such as resistance to corrosion and fire protection, have not received equal attention.

Maintaining existing structures can be extremely costly and therefore, decisions should be taken only after a rational examination of important factors. Structures need to be preserved in order to retard decay; typical activities are cleaning, painting, maintenance of drainage systems, and clearing of expansion joints. Also, structures require repair after damage has occurred. Prior to repair it is important to identify the causes of damage and evaluate whether repair is required immediately. Whereas the search for causes is common practice in many fields, existing structures are often repaired according to the characteristics of the symptom alone. Clearly, accurate knowledge of cause contributes to a successful repair.

Modifications to existing structures should be carried out in such a way that structural integrity is not decreased. Although this criterion appears obvious, there are many cases where structures have been inadvertently weakened by modifications carried out during the life of the structure. For example since 1940, hundreds, perhaps thousands, of riveted bridges have been weakened by modifications employing welding. Ironically, some of these modifications were carried out in order to reinforce the structure. A great deal of specialized knowledge is required to modify existing structures. In Figure 1, two motivations, a new use for the structure and a need to increase capacity, are provided as examples of reasons to modify structures.

Each activity in Figure 1 requires diagnostic or classification procedures to be most effective. These procedures are important for identifying good solutions and areas where more information would be most helpful. Nevertheless, a distinct focus is required for each activity since the user wishes to proceed differently for each case. Therefore, each activity has a unique set of rules which make up and control the methods employed during solution formulation. However, much of the information used by these methods is similar. Also, solutions implemented during different activities can affect each other. Common information requirements and possible interaction can be well accommodated by an integrated system. A proposal for such a system is presented in Section 4.

Many opportunities for creating knowledge-processing systems exist, and work in progress represents a small proportion of possible systems. For any effort, in system creation, a prerequisite for good solutions is a complete definition of the problem. Often, the original definition is inaccurate because domain knowledge as well as the needs of the user have not been represented appropriately. Therefore, an attempt should be made to develop a small prototype as soon as possible in order to begin testing the system at an early stage. An example of such a system is presented next.

3. DEVELOPMENT OF A SMALL SYSTEM

A system called CRACK CONTROL was developed to help engineers decide what to do if a crack is discovered in a steel structure. Intuitive repair solutions such as filling the crack with weld metal do not have the desired effects. Good decisions require a combination of scientific knowledge and experience gained through examining cracks in structures. Generally, if a crack is found in a steel structure, more careful inspection will reveal additional cracks in similar elements. If no action is taken to eliminate the cause of cracking, more cracks usually appear at other locations. These heuristics have an influence upon the knowledge structure described below.

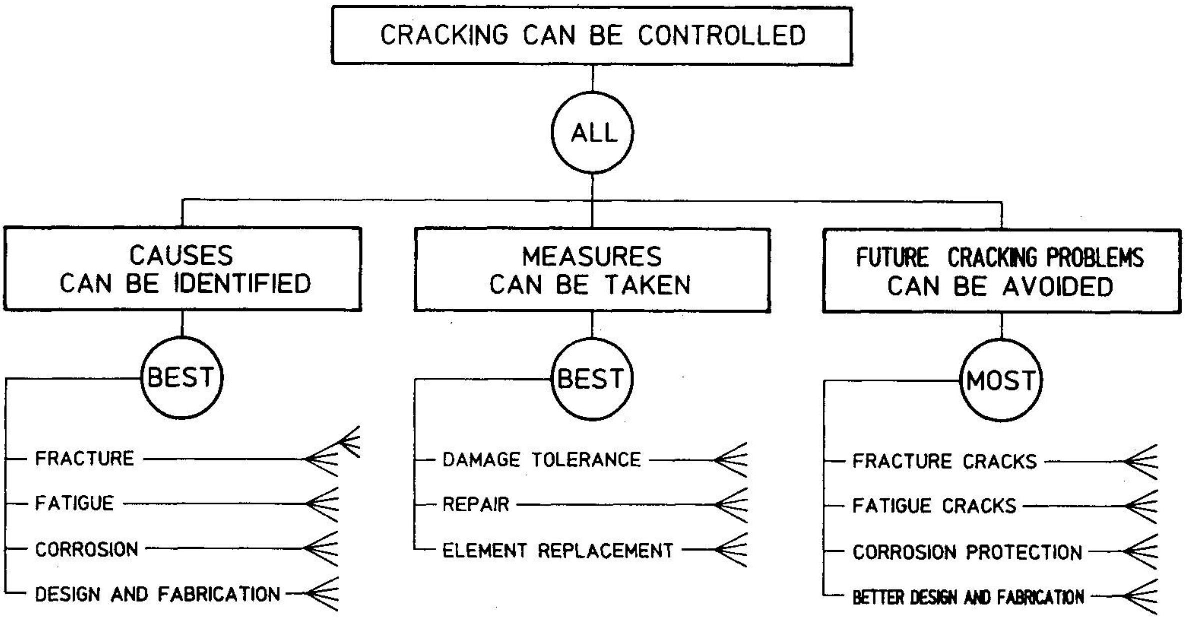

FIGURE 2 *Partial inference net of* CRACK CONTROL

The knowledge necessary to solve this problem is split into three parts, as shown in Figure 2. The first part concentrates on parameters which cause cracks in steel structures and thus, it contains the majority of the diagnostic knowledge in the system. This knowledge is split into categories which reflect the origins of cracking. Cracking may be due to fracture, fatigue, corrosion or design and fabrication practices, or most often, a combination of these factors.

The second part of the knowledge focuses on the most appropriate action, given a cracked element. Measures to be taken are subdivided into damage tolerance, repair and element replacement. Damage tolerance involves no immediate repair but an increased inspection effort. This solution is only explored under certain conditions since it is not appropriate if, for example, further crack growth could cause catastrophic collapse. Repair measures are dependent upon the causes determined in the first part. Element replacement is a valid measure when damage tolerance and repair are not practicable.

The third part of the knowledge concentrates on identifying a maintenance strategy for the rest of the structure. Once cracking has been discovered in a steel structure, the maintenance effort needs to be modified since more cracking is likely. While these considerations do not depend greatly upon the measures chosen for the cracked element, they are closely linked to the causes determined in the first part. Also, several general precautions are needed regardless of the cause of cracking.

This knowledge was implemented rapidly into a small system using a development tool specifically designed for diagnostic applications - THE DECIDING FACTOR (TDF) [6]. This tool was developed using experience gained during the PROSPECTOR project [7] and it has already been employed for diagnostic applications in civil engineering, e.g. [8].

Rather than require direct input of production rules, TDF processes knowledge organized in inference nets, see, for example, Figure 2. The user expresses opinions related to ideas low down on the net. These opinions are transferred into a belief value and multiplied by a factor to contribute to the hypothesis represented as the parent of a set of ideas. In turn, sub-hypotheses contribute similarly to hypotheses further up on the net. Belief values are combined using special logical relationships provided by TDF. In Figure 2, ALL, BEST and MOST are three of eight possible relationships. ALL and MOST pass weighted averages of belief values, whereas BEST passes the highest belief value. Thus, BEST is analogous to OR logic. The system, CRACK CONTROL, employs six relationships in all.

One of the strong points of TDF is the user interface, see Figure 3. Typically, a question screen is composed of an introductory explanation, a question, an answer box and a scale of possible answers. The user manipulates the cursor in order to adjust his answer. He need not reply definitely yes or no. Intermediate answers such as MAYBE SO and THINK NOT are possible. The middle of the scale is the reply, DON'T KNOW. This feature is very useful for applications to structures in service since information is rarely complete and never certain. This interface has been well accepted by users during tests.

Questioning proceeds from left to right in the inference net (Figure 2). It is possible to fix a range of answers, thereby allowing continued investigation of the ideas which contribute to the current hypothesis. If the user replies outside this range, questioning relating to the current hypothesis is terminated, and the system goes on to the next part of the net. For example if damage tolerance was the current hypothesis and the user had any doubt whether further cracking would lead to catastrophic failure, the system would not pursue this possibility further. Therefore, questions relating to the safety and economy of a damage tolerance philosophy would not be asked, and repair would be investigated.

A final step in the system involves a review of the recommendations provided for the particular case. Note that heuristic information is used only to identify the most appropriate recommendations. Once these are identified, the user is asked to what extent he believes that the recommendations can be carried out.

This belief determines which recommendations are reviewed and ultimately used by the system to evaluate the hypothesis that cracking can be controlled.

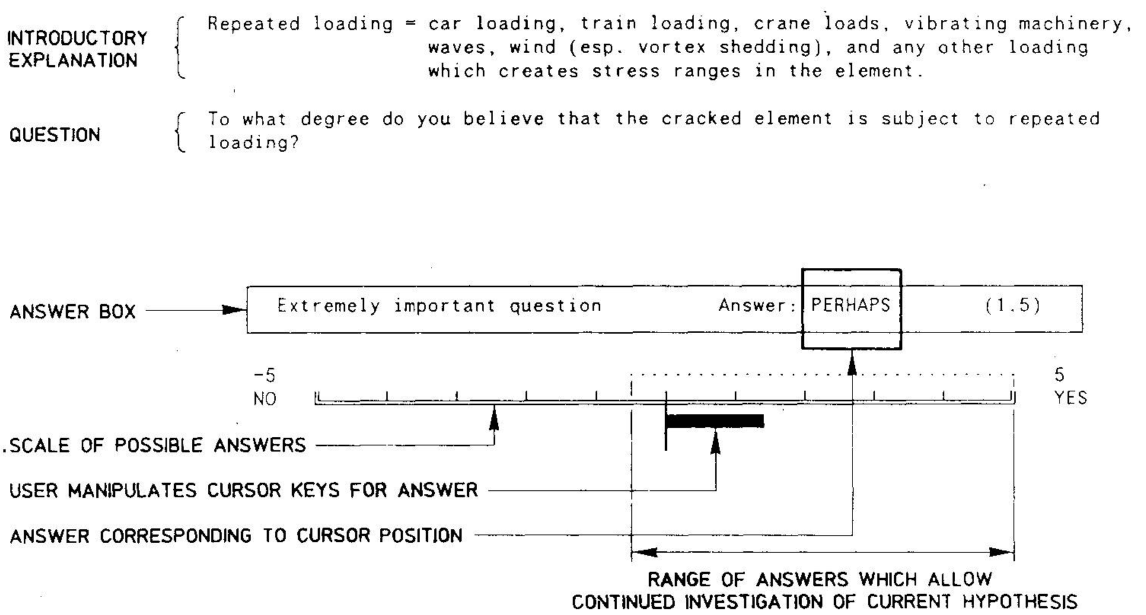

FIGURE 3 CRACK CONTROL *User interface*

Due tô the ease of development, a working prototype was ready for testing two weeks after development began. Many changes were introduced after initial tests. Indeed, it was discovered that the problem was not completely defined from the start. Some measures for dealing with cracked structures were overlooked. Users employ a different language than experts and sometimes prefer that questions are raised in a different order. A small system developed rapidly using a simple tool created a situation where these differences were identified as quickly as possible.

4. HYBRID SYSTEMS FOR LARGE APPLICATIONS

Small systems developed rapidly for testing help to ensure that effort is not wasted solving the wrong problem. Knowledge is verified at an early stage and the requirements of the user become well defined. However, as the size of the problem grows, the number of assertions increases rapidly. Interaction between these assertions becomes difficult to manage and verification of all possible solutions is increasingly arduous. Well organized knowledge becomes essential.

Models and so called deep reasoning provide effective ways to organize knowledge. Generally, two types of models could be used to simulate structures in service. The first type is a mathematical description of the behaviour of the structure. Examples of models of this type include structural-analysis algorithms, fracture-mechanics simulations and fatigue-damage-accumulation techniques.

The second type is a representation where the design and function of the structure is described. Figure 4 gives an outline of such a model of a structure. In this figure, actions, such as gravity loads and wind, act on the structure. The structure is described in terms of the material employed, elements and their connections to each other, details at connections and attachments, built-in stresses, etc. The structure acts on the foundations, which for the purposes of this outline, include surrounding soil and geological properties. External factors, such as salt-water exposure, atmospheric pollution and changing ground-water levels also act on the structure and foundations. Also, changes in the behaviour of the foundations over time may in turn affect the behaviour of the structure.

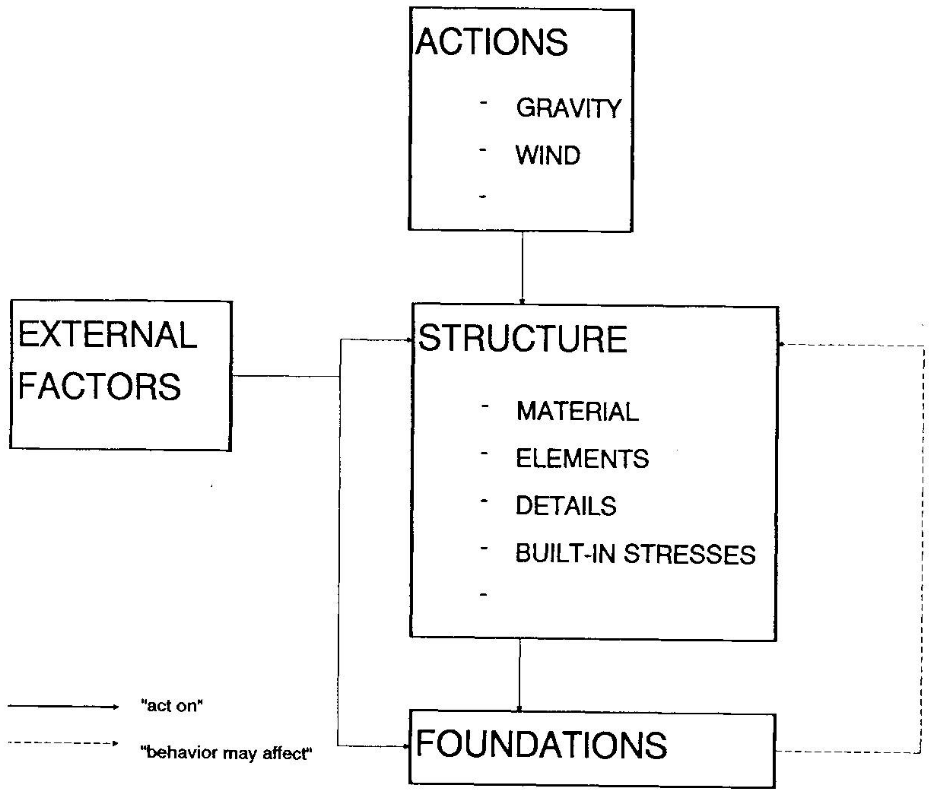

FIGURE 4 *An outline of a functional model of a structure*

Recent work in artificial intelligence has examined the advantages of deep reasoning for diagnostic activities, e.g. [9]. Using models such as the one outlined in Figure 4, domain-independent theories provide methods for diagnosis from first principles. Given a state which is observed to be outside the limits of expected behaviour, models can rapidly identify the origin of faults. They provide a means of representing knowledge for large quantities of information and complicated relationships. Therefore, models are important to the future of large diagnostic systems [10].

A further advantage of models is that they are useful for a range of activities. For example, the model in Figure 4 could assist reasoning during many of the activities shown in Figure 1. On the other hand, shallow systems using heuristic pattern matching are typically constructed to do a specific task.

However, first-principle models [9] are not useful for many types of practical problems. An exact model of the system is required, and uncertain information cannot be treated. As the number of possible faults increases, computational overhead rises exponentially. If multiple faults are considered, models are especially sensitive to problem size. Therefore, first-principle diagnostic models are most useful for medium sized closed-world problems such as small electrical circuits.

Problems associated with structures in service are very different from small electrical circuits. Important information may have a high degree of uncertainty. Relationships between objects may be poorly defined. A structure may have thousands of elements and details, and tens of load cases. In addition, critical measurements may be very difficult to carry out and external factors may include social and political considerations. These factors mean that structures in service have open-world characteristics.

Research into artificial intelligence has developed new techniques which are very useful for representing activities associated with structures in service. For example, specialized strategies used with inexact models may help reduce the effects of the open-world characteristics of existing structures. Rather than attempting to construct complete models, inexact models contain only knowledge relevant to a group of activities [11]. Other developments in non-monotonic reasoning and machine learning have created many opportunities for applications involving ill-defined problems such as those typical of structures in service. These techniques are often implemented within a system which employs shallow and deep reasoning methods.

Until recently, such hybrid systems could only be run on powerful mainframe and stand-alone machines specially developed for symbolic computation. This hardware is not compatible to the needs of civil engineering, and in particular, activities associated with structures in service. Hardware should be portable so that consultation can occur on site. Software should not be dedicated to one machine since several consultations may be needed at different places simultaneously. In the past few years these difficulties have been overcome, thereby creating the conditions necessary for practical applications of large hybrid applications in civil engineering. A summary of the considerations leading to a hybrid approach for large systems in civil engineering is given in Figure 5.

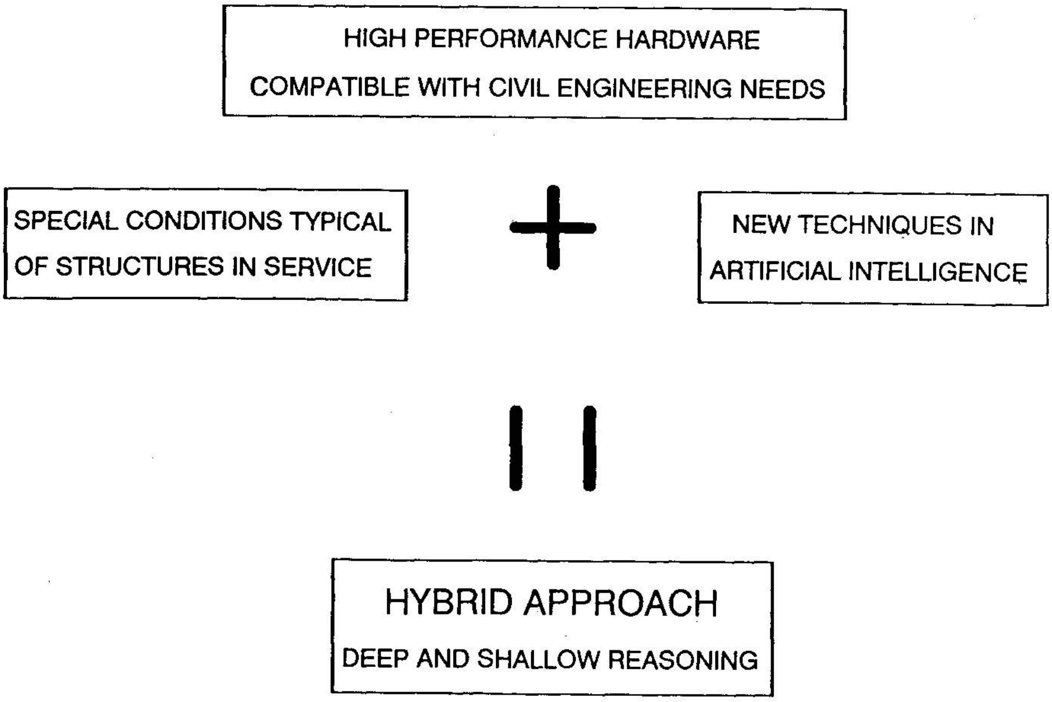

FIGURE 5 *Considerations leading to a hybrid approach for structures in service*

A hybrid approach for activities associated with structures in service is proposed, see Figure 6. The user would start the system by providing information which identifies modules that are appropriate to the problem. The majority of these modules would be activity-dependent. However, some modules, such as those used to estimate behaviour, would be used for several activities. For example, modules such as CRACK CONTROL would be chosen from a library of available small systems. At this point, the system would carry out shallowing réasoning using heuristic knowledge which is independent of the structure in question.

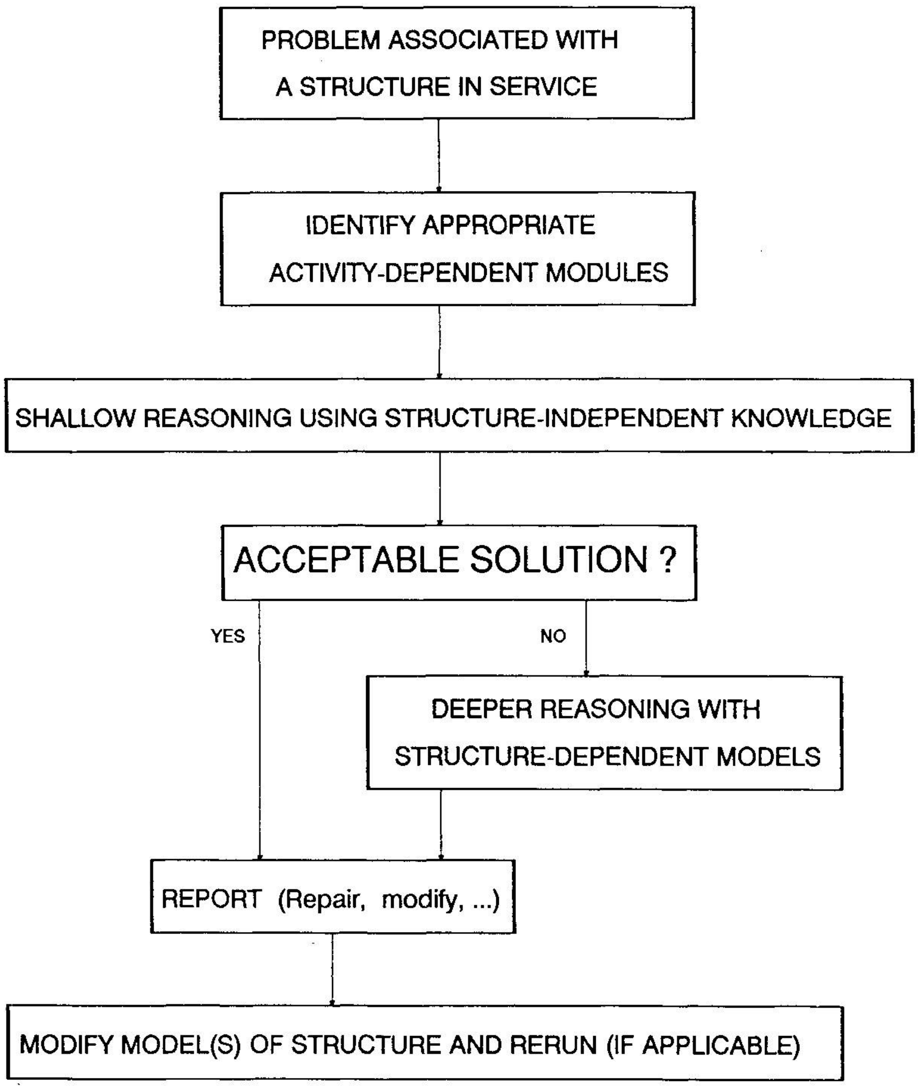

FIGURE 6 *A hybrid-system approach for activities associated with structures in service*

The findings of the system would then be assessed by the user. If an acceptable solution was identified, the system would not invoke methods of deeper reasoning. This step is comparable to traditional engineering methods since engineers typically employ more sophisticated methods when acceptable solutions are unavailable through simpler approaches. Also, if models of the structure do not exist, this step provides an opportunity for evaluating the advantages of creating them. The complexity of some structures in service could require a substantial investment in order to produce useful models.

If an acceptable solution is not identified, the system would employ deeper reasoning using structure-dependent models and more abstract heuristics. For example, if a crack is discovered in a steel structure, reasoning could help identify candidate causes of the cracking by backtracking and examining all factors which affect the element. Optimal locations for additional measurements could be identified and when new information is received, the candidate list would be updated. Most likely causes, learned from previous experience with this structure and others like it, could be placed in default slots; reasoning with such information would proceed until evidence disqualified the assumption. Similar procedures could be employed for identifying other areas at risk in the structure and for evaluation of repairs. As stated already, new research in artificial intelligence has created conditions where these capabilities are applicable to activities associated with structures in service.

The models used would be independent of activities such as those in Figure 1. In this way, information would be shared as required by the particular task. However, many heuristics would be activity dependent, especially those which control how the model is examined. Also, information obtained in the shallow-reasoning stage would be used for pruning search.

The next step is a reporting stage where findings such as recommendations for repair and inspection priorities could be presented. A final step is necessary if any repairs or modifications are carried out. In such cases, relevant models should be revised and the system rerun in order to confirm the success of the changes. In addition, this last step ensures that the models are kept up to date.

Some aspects of this approach are comparable to a multi-level approach developed for fatigue and fracture in bridges [12]. Although this study concentrated on coupling numerical and symbolic computing, many of the advantages of hybrid approaches for different reasoning techniques are demonstrated. Indeed, this study and the approach proposed in this paper suggest that structures in service can be managed more effectively with the help of modern hybrid systems.

CONCLUSIONS

1. There are many possibilities for applications of knowledge-processing technology to activities associated with structures in service. New and current work should improve capabilities to manage structures, thereby reducing costly repairs and unnecessary replacement.

2. Since structures in service experience problems which are almost impossible to define correctly from the start, it is essential that knowledge-base development begins with a rapidly developed prototype for testing with the expert and the user.

3. Models help organize the knowledge necessary for large diagnostic systems. However, for open-world problems encountered by structures in service, a purely model-based system, controlled by domain-independent heuristics, is not appropriate.

4. A hybrid system which combines shallow reasoning with model-based deep reasoning is a feasible and effective approach for structures in service.

ACKNOWLEDGEMENTS

Research at ICOM in the area of knowledge processing is funded by the Swiss National Science Foundation. The author is grateful to N.E. Reed for helpful discussions. Also, the staff at ICOM are thanked for their help in the preparation of this document.

REFERENCES

[1] BAJPAI, A. and MARCZEWSKI, R., CHARLEY : An expert system for diagnostics of manufacturing equipment. Innovative Applications of Artificial Intelligence. AAAI, 1989, pp 178-185.

[2] CHEN S.S., WILSON, J.L. and MIKROUDIS, G.K., Knowledge-based expert systems in civil engineering at Lehigh University. ATLSS Report No. 87-01/1987, Lehigh University.

[3] Expert systems in civil engineering, Proceedings ASCE, 1986.

[4] MAHER, M.L. ed, Expert systems for civil engineers. TCCP, ASCE, 1987.

[5] ADELI, H. ed, Expert systems in construction and structural engineering. Chapman and Hall, 1988

[6] CAMPBELL, A.N. and FITZGERRELL, S. T, The Deciding Factor. User's Manual, Channelmark Corp. 1985.

[7] CAMPBELL, A.N., HOLLISTER, V.F., DUDA, R.O. and HART, P.E., Recognition of a hidden mineral deposit by an artificial intelligence program. Science, 217, 1982.

[8] LEVITT, R.E., Howsafe: A microcomputer-based expert system to evaluate the safety of a construction firm. In : [3], pp 55-66.

[9] REITER, R., A theory of diagnosis from first principles. Artificial Intelligence No. 32/1987, pp 57-95.

[10] MILNE, R., Strategies for diagnosis. IEEE Transactions on systems, man and cybernetics, 17, 1987, pp 333-339.

[11] REED, N.E., STUCK, E.R. and MOEN, J.B., Specialized strategies : An alternative to first principles in diagnostic problem solving. Proceedings of AAAI-88, St. Paul MN, 1988, pp 364-368.

[12] RODDIS, W.M.K., and CONNOR, J., Qualitative/quantitative reasoning for fatigue and fracture in bridges. Coupling symbolic and numerical computing in expert systems, North Holland, Amsterdam, 1988, pp 249-257.

Leere Seite
Blank page
Page vide