Time-critical problem-solving with cached knowledge: a case study in shortwave radio resource allocation

N. Chatterjee[a,*], J.A. Campbell[b]

[a]*Department of Mathematics, Indian Institute of Technology, Hauz Khas, New Delhi, 110016, India*
[b]*Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, UK*

Accepted 1 January 1999

Abstract

Successful communication between two places through shortwave radio depends on the right choice (sometimes within a very short period) of both frequency and power. This choice involves calling on mixtures of precise models and heuristics about prevailing geophysical conditions, and knowledge of specific past experiences. The design and the features of a knowledge-based system that is capable of answering these requirements for a given transmission request are described. There are two unusual constituents: a *caching structure* that can accommodate different knowledge representations and time limits within which answers must be completed, and *interpolation* between items of knowledge in different cells of a cache. Results are demonstrated for simple but realistic radio propagation exercises, to illustrate the effectiveness of caching and interpolation for time-critical knowledge-based computations.

Keywords: Time-critical computing; Caches; Knowledge representation; Knowledge interpolation; Shortwave radio

1. Introduction

This article gives an answer to the question of how to compute with bodies of knowledge, especially when different versions (e.g., rules, cases) of the knowledge are available, in time-critical applications. It makes its recommendations (for how to collect and organise knowledge) and conclusions specific by considering in some detail one application: shortwave radio propagation.

The users of shortwave radio may range from small-scale private organisations to much larger ones such as international broadcasting stations. Also, the users may have their own constraints on transmitter power, usage of certain frequencies, duration of the communication, temporal deadline within which the message has to be communicated, and expected quality of the transmission. For illustration, requirements for *broadcasting* may be clear, interference-free transmission for a considerable duration, whereas an army's priorities may lie in immediacy and not in quality. Such time-critical transmission problems are referred to here as *pressing* ones. For both types of transmissions, an expert's task is to provide guidance in:

- choosing an optimal frequency;
- determining the required power of the transmitter; and also,
- considering whether direct transmission can be achieved, or whether some indirect means should be used.

These decisions depend upon various geophysical properties. The best available guide in this field is still in the form of books (Bird, 1998; Braun, 1982) that provide knowledge about the underlying physics governing propagation of waves between two locations, and/or rough guidance towards selection of ideal frequency ranges; and no straightforward formula or algorithm exists to determine the right choice. Hence an expert depends significantly on past experiences and heuristics to find a solution to a given problem.

The only significant development effort for software in this area that could be identified is from the CALMA project (van Benthem, 1995), jointly commissioned by the Ministries of Defence of France, Great Britain and the Netherlands. However, the aim has been to assign frequencies among the constituents of a network so that their constraints are satisfied and a minimum amount of the electromagnetic spectrum is used. It is therefore different from the intended generalised goal of the work described below.

The present article describes a scheme for capturing an expert's knowledge and then using it in symbolic computations for finding a quick (i.e., adequate to deal even with 'pressing' problems) yet efficient solution for a given problem. The description is intended to illustrate how the knowledge base can be created and structured when one is building time-critical systems in other applications. The next section outlines the domain of shortwave radio propagation. The sections that follow it deal with the structure and use of the knowledge base. Finally the effectiveness of the overall system in solving problems by using different knowledge representations and different accompanying methods of reasoning is illustrated.

2. Shortwave radio propagation: an overview

The range of frequencies from 2.3 to 30 MHz is popularly known as 'shortwave'. For convenience of reference, the part of that range available for broadcasting is divided into several bands: *band 2* runs from 2.3 to 2.495 MHz, *band 3* from 3.2 to 3.4 MHz, and the frequencies between 25.74 to 26.04 MHz comprise the highest band, *band 25*. The propagation properties of signals vary according to the band used. For example, in a simple view of propagation, a signal is radiated from a transmitting antenna towards the sky, and is received at a different location by either a single reflection from the ionosphere or by multiple reflections between the ionosphere and the Earth's surface. The reflection of a signal in the ionosphere depends on two properties: the *strength of ionisation* and the *height of the reflecting level of the ionosphere*. During the course of a day both the height and strength of the reflecting layer vary, causing changes in propagation characteristics. Furthermore, many other geophysical, geographical and/or natural phenomena influence shortwave propagation (Braun, 1982), such as:

- **Daylight path**: The timing when the transmission made—whether it is during daylight hours or at night, or close to a sunrise or sunset period (when what is called the *grey-line effect* in radio engineering jargon occurs).
- **Seasonal effects**: As the solar radiation effects change with the season, the usable frequencies also vary with the time of the year. In particular, one needs to distinguish at least between two seasons, say *winter* and *summer*.
- **Water path**: The reflectivity of water differs from that of land. Waves of higher frequencies can be reflected more easily from water surfaces than from land. Particular attention needs to be given to whether there are long stretches of water in the path (relevant for the case of multiple reflection between the Earth's surface and ionosphere).
- **Sunspot number**: The nuclear activity occurring in the sun, producing light and heat, does not follow a constant pattern. Instead, it is in a state of continual change. However, it has been observed that there are cyclical behaviours of this activity, whose best-known component is called the *sunspot cycle* (Bird, 1998). This is an 11-year cycle during which the intensity of the solar activity increases and decreases in a qualitatively predictable fashion. When the activity is low the ionosphere is less dense overall, and propagation of frequencies higher than 15 MHz by reflection in the ionosphere is virtually impossible. On the other hand, when the sunspot number is high, there will be stronger radiations from the sun, causing higher ion densities in the ionosphere. Consequently, radio waves of higher frequencies can be propagated.
- **Atmospheric noise**: Natural phenomena such as storms, heavy rain, snowfall, and sandstorms release copious electrostatic discharges in the atmosphere, causing radio noise. This noise interferes with the electromagnetic radio wave and may hinder its propagation, besides being heard as an interfering distraction by the radio listener.
- **Auroral activities**: Electromagnetic discharges, primarily caused by flows of ions from the Sun, into the magnetic polar regions, cause interference in the propagation of radio waves. This is most acute for any long-distance propagation, especially when signals may have to cross an auroral zone once or even twice on their path.

Additionally, there are some human-created obstacles, such as:

- **Undisciplined transmissions**: Many local radio centres often do not abide by the international rules, and broadcast programmes on certain frequencies (aiming at some local population) which block the transmissions that are officially assigned there.
- **Jamming**, i.e., deliberate interference from certain

countries, primarily for political reasons: These countries broadcast various loud noises, sometimes interspersed with real or imitation coded messages (e.g., old programme tapes played backwards) in order to block the broadcasts from other sources in some target areas.

An expert in shortwave propagation needs to consider all these factors, and others, in order to determine a good frequency and power for a given transmission request. Further knowledge is required for suggesting indirect transmission should the prevailing conditions be found unsuitable for a direct one. Indirect transmission may be of various types such as using a distant transmitter connected to the source via a telephone line; or using an off-air relay. The relevant knowledge exists in a mixture of forms, including existing procedural programs, rules and cases. Their use is further complicated by the question of which one to choose when a request is time-critical.

One way to deal with time-critical requests to heterogeneous knowledge bases is to use some form of *cache*. A suitable caching scheme is described in the next section. After that, the three major aspects of building this kind of knowledge-based system are considered:

- identification and acquisition of relevant knowledge;
- representation of the knowledge in the system;
- designing an efficient reasoning scheme to exploit the acquired knowledge in deriving solutions.

3. Fundamentals of the cache-based scheme

The radio domain is one of many in which experts do not stick to a single reasoning mode in solving problems. The method that they adopt (perhaps without conscious realisation) depends on the nature of the problem. For example, if the problem demands a quick solution, an expert often prescribes some answer (here, a frequency and a power) from some distillation of prior experience. On the other hand, when more time is available, a radio expert examines various geophysical factors currently prevailing, and works out their influences to check whether a direct transmission is possible or not. If it is possible, the expert modifies the initial intuitive values according to a perception of the special prevailing conditions; otherwise an alternative method of transmission (e.g., relay or use of some other transmitter) is suggested to carry out the task. There are, however, situations when due to some peculiarities in the transmission conditions even such an elaborate solution fails. In these situations an expert usually resorts to specific past experiences where similar problems have been faced, and designs solutions to imitate these past similar situations. The cache-based scheme has been designed in order to accommodate this multi-paradigm style of reasoning.

Knowledge representation here is in the form of a multi-level cache. Structurally the cache can be pictured as a multi-dimensional ($1+k$, say) array. The first dimension is called the *level*, and the other k dimensions are *column dimensions* of the cache.

Each level is associated with some type of solution or solution scheme. Which particular level of the cache will be accessed while solving a given problem depends upon the available computational time. For a *pressing* problem, this time-limit T is an input, and the level accessed is the one for which the expected time for the resulting computation is no greater than T.

Each of the k column dimensions (A_i, $i=1, \ldots, k$) refers to some key feature or abstraction that is essential in describing a problem situation. Each of these dimensions is then represented by a number of columns each representing some particular values (or ranges of values) V of A_i. The j-th of these values for A_i is denoted as V_{ij}. The j-th column holds knowledge about how to solve problems better characterised by the particular value V_{ij} of A_i than any other V_{ih} for $h \neq j$. Fig. 1 shows a cache with four levels, and two column dimensions. The number of V values corresponding to A_1 is five, while for A_2 it is three. A solution for a given problem may now be sought from the cache via a slice through the cache identified by $(1+k)$-tuples $(x, v_1, v_2, \ldots, v_k)$, where v_j $(j=1, \ldots, k)$ are the problem characteristic values for the k column dimensions of the cache, and x is a level index. This slice reduces to a single cell when a particular value of x is chosen, and each v_i (the problem value for the ith property A_i), matches exactly with one of the V_{ij}s. When some v_i does not match exactly with any of the V_{ij}s, one can use some metric to find the distances of the V-values from the given v, and select cache cells on that basis.

For this scheme to work, it must be possible to identify an order on the possible values of each A_i $(i=1, \ldots, k)$, and then to set up a metric in order to assign distances between any two columns. This metric and the ordering imposed along each column dimension allow a position to be found for a given problem in that dimension. If the values are symbolic, this order is not immediately obvious. But it is usually not difficult to impose one if only a single property that is important for the solution of problems of the underlying symbols is considered. For example, there is no obvious order on the animals deer, cow and elephant. But one can hope to find some order with respect to a suitable *specific* property: cow falls between deer and

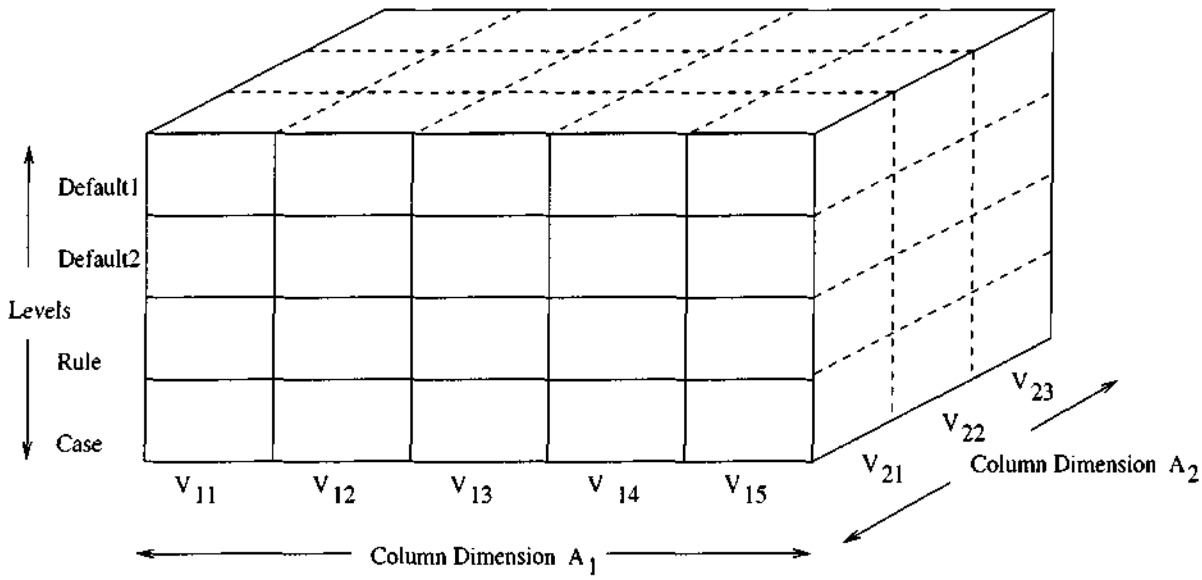

Fig. 1. Schematic diagram of a cache.

elephant when the most significant property is weight, when speed is important, elephant comes between cow and deer, while deer falls between elephant and cow under 'potential for domestication'. More discussion on this issue may be found in (Chatterjee and Campbell, 1994). Section 4 covers selection of abstractions (as column dimensions) and determination of their relevant orders for the radio cache.

As successively deeper levels of a cache will contain successively more thorough (thus requiring more time for computation) ways of answering requests, the form of the knowledge held in different levels will usually vary as well. The radio cache has been built with the following contents:

- *default solutions*, i.e., intuitive ready-to-use suggestions;
- *rules*, i.e., domain-related rules containing knowledge about how to improve upon default solutions taking prevailing transmission conditions into consideration. If the number of rules corresponding to a cache cell is too large to be held in the cell, the rules may be stored in a separate rule base and pointers may be used from each cache cell to the relevant rules.
- *cases*, or past experiences of the experts that may be used to deal with special or unusual problems in transmission. For cases too, pointers may be maintained from the cache cells to appropriate entries of the system's case base.

The key features/abstractions that make the column dimensions of the radio cache are: the location of the transmitter, the target area, the transmission time, and the season when a transmission takes place. These lead to five column dimensions for the radio cache, as explained in Section 4.1. That section may also be viewed as a case study for how one should set up a cache scheme for other applications.

3.1. Reasoning with the caching scheme

The scheme is provided with two inputs: the *problem description* and the *temporal bound*. Reasoning with a cache-based system essentially consists of four major steps.

3.1.1. Identification of the right cache level

The *i*th level of an *m*-level cache is associated with a characteristic computing time T_i, which is a conservative estimate of the time needed to arrive at a solution by using that level of the cache. Given a time T by which a solution must be computed, the best level x is chosen so that $T_{x-1} < T < T_x$, for $1 \leqslant x \leqslant m$, where T_0 is the minimum time that has to be allowed to derive the default solution. If T is smaller than T_0, the software using the cache should refuse to provide any answer, and should report that insufficient time is available. This is because the system needs some overhead time for preliminary analysis of a situation and for access to the cache. For practical applications one may assume that the time bound T will be greater than T_0. The T_is can be static, for example, derived from previous experiments, or dynamic, with updates determined during actual use.

3.1.2. Determination of the right cache cell(s)

As indicated earlier, at any level of the cache each cell can be regarded as the intersection of k columns, when there are k column dimensions. The most suitable cache cell for a given problem can be determined by finding the V value, in each dimension, that is closest to the problem's v value (defined in Section 3) in that dimension. This specifies the cell that lies at the right intersection for the purposes at hand.

It is, however, unrealistic to assume that an incoming problem will strictly match a single V in each column dimension so that only one cache cell will be relevant for the answer. In general, it is likely that any v will fall between two V values, so that two columns in one dimension may appear to be significant. Thus, for a cache having k column dimensions, one may find that as many as 2^k cache cells are selected. Finding a solution by using the contents of so many cells is not advisable. The number of cells is pruned in the following way.

When in some column dimension a problem's v matches exactly with a V, the associated column is selected with weight 1.0. When it does not, two columns corresponding to the two closest Vs are considered, with appropriate weights. The weight assigned to a column is in inverse proportion to the distance (according to some chosen metric, see Section 3) of its V values from the relevant v of the problem feature. The weight of a cell is defined to be the product of the weights of its constituent columns. The contents of the two cells with the highest weights are then used for computing the final solution.

3.1.3. Retrieval from the cache

The prescribed cell(s) are accessed to retrieve the information required for a solution. The retrieved may be stored as a solution (as in the 'default' level), or as pointers to an appropriate section of the knowledge base (as in 'rule' and 'case' levels) from where relevant items are retrieved. The retrieved information is then used in determining a solution for the current problem.

3.1.4. Framing the actual solution

When information is retrieved from two cache cells, the reasoning scheme needs to consider the contents of both the cells for computing any solution. Any such attempt requires some means of balancing the two suggestions: in effect, doing something analogous to interpolation in numerical analysis using the weights computed in the second step above. Section 5 illustrates the use of such interpolations in framing solutions from the contents of two cells.

4. Knowledge acquisition for building a cache

Building an actual cache requires identification and representation of the various fundamental or abstract domain features for the relevant application. In practice, one should distinguish among four classes:

- **Problem-related**: features that are required to register the essentials of a problem. These features are used to determine the column dimensions of the cache.
- **Retrieval-related**: features that are exploited in retrieval of appropriate rules and cases from a knowledge base.
- **Solution-related**: knowledge about the natures of the solutions (such as the permissible actions and their parameters).
- **Auxiliary**: other pieces of knowledge that help in decision-making.

Fig. 2 shows the roles of the various types of knowledge. The next section discusses details of all these types of knowledge for building a cache-based system.

4.1. Identification of problem-related features

These >are the features that together convey the nature of the problem to the system, and ensure that the cache columns and cells are labelled accurately.

For the radio domain, seven features are important for gross characterisation of a problem: latitude and longitude of the source, the same quantities for the target, information about the season, and the beginning and end times of the transmission. However, in the actual application it is possible to compress this information into five variables, giving five column dimensions for the radio cache. They are: **source-latitude**, **target-latitude**, **longitudinal-difference** (i.e., the difference in the longitude values for the source and the target), **seasonal effect** (which takes care of seasonal variations), and **daylight-path** (the condition of daylight and/or darkness along the propagation path).

The V values for the columns in the five dimensions are as follows:

- For the two latitude-related dimensions (**source-latitude** and **target-latitude**) it is enough to use five approximate positions on the globe. They are: *N-arc* (Northern Arctic), *N-temp* (Northern Temperate), *Tropical* (Equatorial region), *S-temp* (Southern Temperate) and *S-arc* (Southern Arctic). If the latitude values for a given location is close to the border between two regions, both the Vs may be selected during retrieval, with each receiving a weight inversely proportional to the distance of the

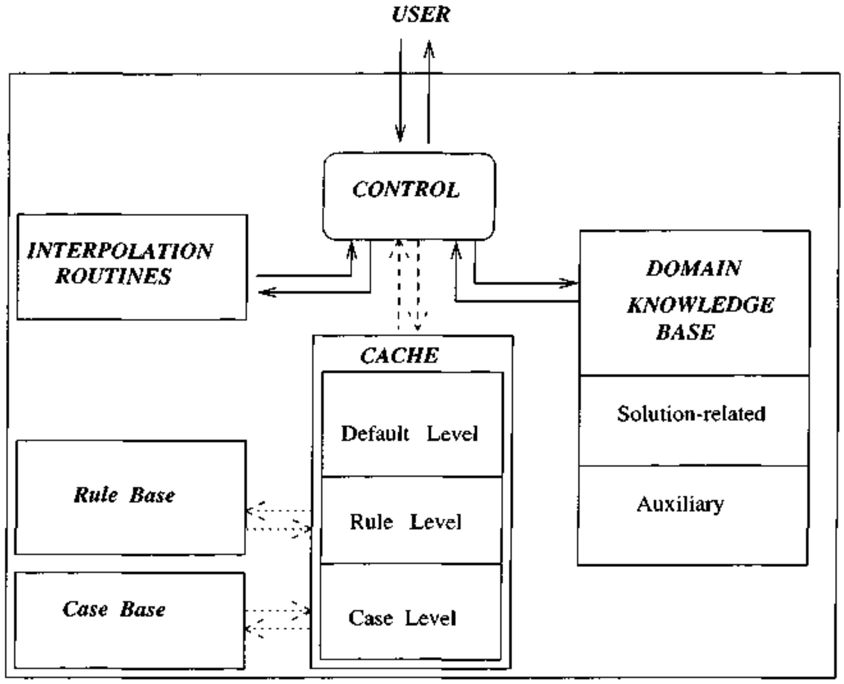

Fig. 2. Schematic diagram of a cache-based system architecture (dashed lines = communications through *problem-related* features; dotted lines = communications through *retrieval-related* features).

given location from the midpoints of the regions specified by the two appropriate *V* values.

- The exact **longitudinal-difference** between two places is not needed for an expert's decision; a qualitative measure of the difference is sufficient. The actual range is mapped into a set of four *V* values: *low* (0–20°), *medium* (21–50°), *high* (51–100°), and *very-high* (101–180°). For a situation close to a border between ranges, the same scheme of weighted selection of two columns that is given above may be used.
- For **seasonal effect**, experts distinguish between just two values: *summer* and *winter*. If the season actually is 'autumn' or 'spring', both columns are used with equal weights.
- For **daylight-path**, too, there are only two basic situations: *day* and *night*. For any given problem, finding the correct value is just a matter of spherical trigonometry. Additionally, if the path is 'grey', which happens when a source or target is close enough to local sunrise or sunset, the two columns with *V* values 'day' and 'night' receive equal weightings.

The example above illustrates the kind of reasoning that constructors of a cache should adopt initially: identification of key features or abstractions that can be used to index relevant knowledge for a field, choice of values for those entities (to label parts of the cache) for partitioning the stock of knowledge in ways that respect the underlying expertise, and specification of tactics to deal with problems that fall near partition boundaries within the cache.

Treatment of the caching of knowledge for short-wave radio propagation continues for the rest of Section 4. Each subsection is both a particular account of characteristics of the chosen application and a demonstration of how to go about capturing knowledge for any application where a cache is appropriate.

4.2. Identification of retrieval-related features

As the nature of the knowledge required for carrying out a reasoning task varies with the cache levels, the retrieval-related features also vary with the levels.

For default-level reasoning no additional knowledge is needed beyond the cached entries. However, reasoning at deeper levels needs some additional knowledge. For example, special treatment may be required for certain particular atmospheric conditions, such as whether the signal is to pass through either of the Polar regions, or whether primarily water cover exists along the propagation path. An expert uses rules and other forms of heuristics to modify the cached answers. For example, one pertinent rule is:

IF water cover of the path is *more or less high*
THEN increase the default power by 10%.

Typically, for any rule-based systems, not all rules are relevant for each particular transmission request. The above rule is required only for situations where the distance between the source and target is very high, so that multiple reflections between the ionosphere and land surface take place before a signal reaches its target. At the rule level in the cache, there are thus pointers to it only from cells where it is relevant.

The same observation is true of cases, for the deepest or case level of the cache. The relevant pointers can be established once one identifies where each available case belongs in terms of the indexing that has been described in Section 4.1.

4.3. Identification of solution-related features

Under this heading, all the tasks that contribute directly to achieving a solution are considered. For the radio propagation problem there are three such tasks:

- *comp-freq*, i.e., computation of frequency;
- *comp-power*, i.e., computation of power; and
- *comp-transmitter*, i.e., deciding which transmitter should be used.

Each of these is treated separately below.

4.3.1. Computation of frequency

First, the frequency band (Section 2) that suits the prevailing conditions is determined. Given the band information, a choice about frequency is made from the alternative frequencies from that band, so that frequencies where interference is likely are avoided. This requires inspection of a suitable database (Magne, 1998) to find whether a certain frequency is likely to be blocked for transmission (due to interference) for the given target location at the given time.

If no suitable frequency is found in the suggested band, the search is extended either to the next band if it is a 'broadcast' problem, or to the frequencies just outside the suggested band if the problem is of 'pressing' (Section 1) type.

4.3.2. Computation of power

Determination of the desirable power to make a transmission effective depends in part upon whether any upper limit exists on the usable power at the transmitting location. If there is such a limit, the use of a relay to bypass that constraint can be suggested, or if an urgent transmission is necessary, use of maximum available power may be recommended (compromising on quality for speed). There are other recommendations, such as *use-suggested-power*, *use-boosted-power*, all involving computations at the rule level to modify the suggested default-level power. At the case level, relevant power information is already contained in the records of past cases.

4.3.3. Selection of a transmitter

Selection here may require computations, involving rules and/or specialised algorithmic knowledge, leading to one of three possible recommendations:

- *use-local-transmitter*, i.e., using the transmitter that is locally available as it is.
- *use-boost-transmitter*, whereby the capacity of the local transmitter is increased by introducing extra transmitter(s).
- *use-relay-transmitter*, the transmission from the source is relayed to the target using a relay-transmission. end

4.4. Identification of auxiliary knowledge

One should regard as 'auxiliary' any knowledge that is accessed through use of cached knowledge, but that is not appropriate to locate in the cache itself. The criterion for appropriateness usually reduces to one of efficiency or economy. For example, there may be items of knowledge that are referenced from random places in the cache, and always in the same form. Part of the knowledge-acquisition process for any application where caching is relevant is the identification (and representation) of the material that does not belong in a cache.

In the particular example of the shortwave radio problem, most of the 'knowledge' in fact amounts to data, such as:

- *latitude-longitude table*, the latitudes and longitudes of cities;
- *blocked-frequency table*, information about which radio station broadcasts at which frequencies at a given time of the day;
- *linkage-availability table*, information about the existence of wired communication links between different locations, such as broadcast sources, or transmitters.

However, knowledgeable access to the data may require a level of interpretation. Heuristics to interpret these data may constitute a set of rules which belongs outside a cache.

The cache-based system for solving the shortwave radio propagation problems has been designed and built according to the above-mentioned considerations. The performance of the system has then been evaluated by experiment. The aims of the evaluation have been:

(a) to establish the viability of a knowledge-representation scheme, namely the cache, that can cater for different reasoning tactics under constraints on computing time, while also using a single uniform method of indexing.

(b) to check that qualitatively better solutions are available as one moves to deeper levels of the cache.

5. Experiments and results

The experiments involve a representative set of problems that a full-fledged system or an expert should be expected to handle. The experiments also cover each of the three types of cached answers: *default*, *rules* and *cases*. For the radio cache the 'default' solution comes as suggestions for the 'lowest' and the 'highest' bands (in MHz) within which the frequency must lie, the 'best' band i.e., the one that should give the best transmission, and 'power', i.e., the desired power (in kW) to carry out the transmission.

Depending upon how reasoning is accomplished on the retrieved cached answers, two levels have been designed for 'default' solutions: Default1 and Default2. At 'Default1' level, in order to answer extremely time-critical demands, only a single cache cell (the one with highest weight) is accessed for framing solutions. In 'Default2' level, two cells are accessed (as explained in Section 3.1), and the result obtained by applying interpolation on their contents is used in designing solutions. In 'rule'-level reasoning, the interpolated result of the 'Default2' scheme is further improved using domain-dependent rules.

The examples below illustrate what happens at different levels of the cache, for three representative problems.

5.1. Results at Default1 level

5.1.1. Example 1

Consider a transmission request from Calcutta to London for the month of June between 1100 and 1145 UTC (a modern equivalent of GMT). The single cache cell that the method chooses is designated by the V values *tropical, northern-temperate, high, summer* and *day* for the five abstractions, 'source-latitude', 'target-latitude', 'longitudinal-difference', 'seasonal effect' and 'daylight-path' (Section 4.1), that characterise the column dimensions for the radio cache. The computed result is:

Transmission from Calcutta to London
Time of transmission (UTC): 1115 to 1145
Seasonal-effect: summer
Daylight-path: day
Use Frequency 15530 kHz
Use Power 110 kW
Use Transmitter local-transmitter

5.1.2. Example 2

Here a transmission request between the same two places as above is considered, but in December between 1835 and 1920. The output is then:

Transmission from Calcutta to London
Time of transmission (UTC): 1835 to 1920
Seasonal-effect: winter
Daylight-path: night
Use Frequency 9430 kHz
Use Power 120 kW
Use Transmitter local-transmitter

5.1.3. Example 3

A request for transmission between Auckland and Johannesburg in December (summer for the southern hemisphere) from 1830 to 1920 UTC yields the result:

Transmission from Auckland to Johannesburg
Time of transmission (UTC): 18.35 to 19.20
Seasonal-effect: summer
Daylight-path: day
Use Frequency 15615 kHz
Use Power 100 kW
Use Transmitter local-transmitter

In real life, Radio New Zealand does not have any programmes specifically for South Africa. However, the best receptions reported in 1995 to radio club magazines at the specified times, under the ionospheric conditions prevailing when this test was made, were in fact on the same band as for the frequencies given in the results above.

5.2. Results at Default2 level

Reasoning at Default2 level considers two cells in order to frame slightly better (in principle) solutions. First the values are retrieved from each of the two cells, and then they are interpolated (using the weights of the selected cells) to produce the final suggestion. For illustration, consider the problem given in Example 1 above. Reasoning at Default2 level chooses the two cache cells indexed by ('tropical', 'N-temperate', 'high', 'summer', 'day') and ('tropical', 'N-arctic', 'high', 'summer', 'day') respectively, with weights 0.52 and 0.48. The interpolation now takes place in the following way:

To compute the frequency, the best bands suggested in the two selected cells are considered. For the given problem these are bands 15 and 17 respectively. These two values are next subjected to weighted interpolation to determine the ideal frequency—or to be more precise, to determine the best band and which region of frequencies should be searched for a unoccupied frequency. The interpolated result is (*17 'lower'*). The lower-end frequencies of band 17 are then searched for an unoccupied frequency, with the result shown below. Computation of power is performed using straightforward weighted numerical interpolation on the recommended powers, i.e., 100 kW and 120 kW.

Thus use of the second cell in decision-making has resulted in a higher frequency and higher power. This is because although London is in the 'Northern Temperate' zone, its latitude value is quite high, and therefore some influence of the considerations for transmitting to the 'Northern Arctic' zone have been involved. The final suggestion is computed as:

Transmission from Calcutta to London
Time of transmission (UTC): 1115 to 1145
Seasonal-effect: summer
Daylight-path: day
Use Frequency 17510 kHz
Use Power 115 kW
Use Transmitter local-transmitter

Similar improvements over the results obtained through reasoning in Default1 level have been obtained for other problems. For illustration, the solution given at this level to the problem stated in Example 3 above reads:

Transmission from Auckland to Johannesburg
Time of transmission (UTC): 1835 to 1920
Seasonal-effect: summer
Daylight-path: day
Use Frequency 15090 kHz
Use Power 110 kW
Use Transmitter local-transmitter

The results obtained from the knowledge-based computations have been assessed, by experts who have examined them, as being of reasonably good quality. For an arbitrary user, to ascertain the quality of the answers, it is suggested that the result obtained from the cache be compared with commercially available tables for selecting suitable bands, as given in (Jacobs, 1992). In these tables the best MHz values for transmission between two large geographical regions (such as Europe or South-East Asia) are tabulated for different time periods of a day (normally expressed as 4-hour intervals) and for the months of the year. However, these tables are not foolproof, and not all the frequencies in the band behave in the same way in supporting a transmission task, as they may be subject to interference by some other broadcast at the same time. But experts' comments on the viability of the frequencies suggested are based on their experience and observations on various past transmission exercises, where further special features are involved, and are therefore less likely to have shortcomings.

5.3. Results at the rule level

Reasoning at the rule level also considers the two selected cells, but unlike the situation for Default2 level reasoning, here the method improves upon the suggested solution by considering rules associated with the cells under consideration.

When the Example 3 problem is treated at the rule level, a relay-based transmission (due to likely polar auroral disturbance of direct transmission) is suggested. The most relevant rule is:

IF problem type is *broadcast* AND polar-influence is *somewhat possible* THEN use relay-based transmission.

The result obtained here is:

Transmission from Auckland to Johannesburg
Time of transmission (UTC): 1835 to 1920
Seasonal-effect: winter
Daylight-path: night
use-relay-transmitter—relay-centre is Colombo
Transmission from Auckland to Colombo
Use Frequency 11630 kHz
Use Power 100 kW
Use Transmitter local-transmitter
Transmission from Colombo to Johannesburg
Use Frequency 11605 kHz
Use Power 60 kW
Use Transmitter local-transmitter

Also, at this level the computation deals with the issue of making a good choice of transmitter (unlike the default levels, where only suggestions regarding frequency and power are made).

For Example 1 and Example 2 the latitude and longitude values of the source and destination suggest that the transmission path does not intersect with the Polar region. Hence the computation adheres to the default plan, i.e., direct transmission. But there are rules that may qualify the solutions for problems where other factors apply. Two experiments involving small variations of the problem specified in Example 2, illustrate this point.

5.3.1. Experiment 1

Suppose that the transmitter under consideration has an upper limit of 100 kW in power, but it is also known that the transmitter centre can be connected through a reliable telephone line to some remote transmitter so that the effective power for the transmission can be boosted. In this arrangement, the computation works as follows:

1. For estimating the required power, the following two rules are the most appropriate:

   IF water cover of the path is *more or less high* THEN increase the default power by 10%.

   and

   IF auroral spread of the path is *between 4° and 15°* THEN increase the default power by 20%.

The method applies the rules on the default values suggested in the selected cells. In this problem, the two selected cells and their weights are computed in the manner discussed in Section 3.1. When applied to the cached value of the first cell (which is 120 kW), the two rules produce suggestions 132 and 144 kW respectively. The same rules, acting on the cached value obtained from the second cell, give 121 kW and 132 kW. Weighted interpolation on these suggestions yields a recommendation of $132.48 \approx 132$ kW.

2. From the assumption that the maximum usable power is 100 kW, the method therefore concludes that a boosted-transmission would be required.
3. The next step is to estimate the required frequency. The relevant rules here are:

   Rule A: IF effect of greyline is *nil* THEN use the suggested *best band*;

   and its complementary rule,

   Rule A^c: IF effect of greyline is *100%* THEN use the suggested *lowest band*;

   plus a third rule:

   Rule B: IF water cover of the path is *more or less high* THEN use the band that is *one band higher than the best.*

Since the transmission period is such that a greyline effect (Section 2) is guaranteed, rules A^c and B apply completely in the given context. The results of their applications to the selected first cell suggests that the best bands are 7 and 10 respectively,[1] while the corresponding results obtained from the second cell are 4 and 8. The weighted interpolation on these values suggests the use of a frequency from the lower end of band 7. The *blocked-frequency table* (Section 4.4) is now searched to obtain a clear frequency for transmission to London, which turns out to be 7110 kHz. The overall output then reads:

Transmission from Calcutta to London
Time of transmission (UTC) 18.35 to 19.20
Seasonal-effect: winter
Daylight-path: grey-line-time
Transmission from Calcutta to London
Use Frequency 7110 kHz
Use Power 130 kW
Use Transmitter boost-transmitter-via-telephone-line

By contrast, for the sake of experimentation, a solution was sought for a similar transmission (i.e., Example 2), by advancing the transmission time by 1 hour (so that the greyline effect becomes 60%). Here the computations found that rules A and A^c received weights of 0.4 and 0.6 respectively, causing the selection of a frequency of 7510 kHz.

5.3.2. Experiment 2

The same problem has been considered here, but with the assumption that the telephone connection in the source location is not sound enough to make a good-quality transmission. This compels the system to abandon the possibility of a direct transmission. The recommendation then obtained is:

Transmission from Calcutta to London
Time of transmission (UTC): 1835 to 1920
Seasonal-effect: winter
Daylight-path: grey-line-time
abandoning direct transmission because not sufficient poweruse-relay-transmitter—relay-centre is Dubai
Transmission from Calcutta to Dubai
Use Frequency 9965 kHz
Use Power 70 kW
Use Transmitter local-transmitter
Transmission from Dubai to London
Use Frequency 9430 kHz
Use Power 105 kW
Use Transmitter local-transmitter

In terms of comparison with actual propagation, these recommendations amount to a qualitative improvement in the solutions. This is a consequence of the domain-specific rules. As indicated earlier, not only is the improvement confined to a better estimation of frequency and power, but it also discounts the feasibility of a direct transmission and readjusts the plan by suggesting an appropriate relay transmission.

[1] Officially, band 10 does not exist: there is no shortwave broadcasting allocation between 10 and 11 MHz. But for the purpose of interpolation it is considered as an interpolating value.

5.4. Reasoning at the case level

The case level gives access to further specialised experience; in particular, in dealing with unacceptable audio quality and problems with transmitter allocation.

Past cases to which an expert will refer do not necessarily describe transmissions between the same source and target as in a current problem. Therefore the solutions proposed there may not be directly applicable. The method refers to past cases to guide only the overall approach. The actual values for frequency and power need to be determined from the context of the current problem, and/or from the cached answers. The overall solving procedure at this level follows four primary steps:

1. Estimation of initial values for the key parameters;
2. Asking the user to indicate shortcomings in the proposed solution;
3. Retrieval of cases that are appropriate in that context;
4. Adaptation of the cases to the current situation.

5.4.1. Estimation of key parameters

At this deepest level of the cache, instead of using the two seasonal *V* values (*winter* and *summer*) only, many intermediate situations such as *almost summer*, *between summer and winter* are considered. For this level, the user is invited to state the season for the current problem in terms of these expressions. The method then interpolates (with this description) on the cached answers for the entries stored with the *summer* and *winter* columns to obtain its output. For example, to carry out daytime transmission from Calcutta to London, the cached values for summer with respect to the 'lowest', 'highest' and 'best' bands (in MHz) and 'power' (in kW) are: 9, 21, 15 and 110 respectively. Values for the same parameters for winter are: 7, 17, 11 and 120. If the transmission season is specified as 'between summer and winter', the method interpolates on the summer and winter values of the same parameters with equal weights. The result thus obtained is: 8, 19, 13 and 115 representing the four respective key-feature values. On the other hand, if the transmission time is specified as 'almost summer', the weights for the two cache cells concerned alter, and the result then obtained is: 8.8, 20.6, 14.6 and 111.1.[2]

5.4.2. Retrieval of appropriate cases

Each cache cell contains pointers to a list of cases that are relevant for the related transmission conditions. Ideally, some of the stored cases should match with the current problem in the source and target as well. However, it is not practicable to collect cases matching every situation. Hence, for a given problem, the geographical distance between its source and the source location of each relevant case is determined. Regarding the target, the system considers those cases where the target is in the same region (e.g. North America, Central Asia, Western Europe) as for the given problem.

The ultimate selection is then made by taking these factors into account, as well as some secondary features whose natures depend on the type of difficulty involved. For example, corresponding to transmitter-related difficulties, relevant features may be:

- upper limit of usable power;
- reliability of telephone line (in the source location);
- list of relay centres (which may be connected to the source).

When a knowledge base does not yet contain information about those phenomena (as in this project), the computation depends on the user's feedback in this regard. Users may avoid exact quantification and use symbolic values such as *high*, *low*, or fuzzier descriptions such as *rather high*. Cases retrieved using these values are then subjected to interpolation. To interpolate on fuzzy terms, the standard scheme (Zadeh, 1979) for deriving fuzzy membership values for different fuzzy hedges is followed with extensions to additional fuzzy terms including *somewhat*, *fairly*, and *rather*. Details of the scheme are presented in (Chatterjee and Campbell, 1996).

For illustration, consider the problem given in Example 1 earlier. The method first suggests a Default2-level solution: frequency 17510 kHz and power 115 kW (as in Section 5.2). Suppose now that this suggestion results in poor audio quality, and the cause has been diagnosed as a *rather high* amount of solar activity. The associated cases retrieved (for 'unacceptable-audio') are:

(a) from Karachi to London, where solar activity has been described as 'high';
(b) from Tehran to London, where the solar activity has been 'low';
(c) from Colombo to Moscow, where the solar activity has been 'very high';
(d) from Kazakhstan to Frankfurt with 'very high' solar-activity.

First, the distance of each case from the current problem is measured. The overall distance of each case is the sum of the absolute differences between the values

[2] The decimal values in the bands do not have any detailed significance but in this application this fraction is useful to determine the frequency range within a band's limits where the search for a frequency that is free of likely interference should be made.

of each relevant parameter—source, target and solar-activity level. The values for all parameters are normalised so that they lie between 0.0 and 1.0. For source and target locations, the geographical distance is divided by the half of the Earth's circumference (i.e., the maximum possible distance between two locations). This kind of metric is standard in numerical taxonomy (Sneath and Sokal, 1973). For solar-activity level, distance is computed using the fuzzy membership values stated above.

The two cases having least distances from the current problem are then selected for computing the final solution. For Example 1, they are the cases (a) and (d) given above. The weight of each case is then normalised, and the weights computed for (a) and (d) are 0.74 and 0.26. The cases are now subjected to 'case interpolation' in order to derive a solution.

5.4.3. Case adaptation through interpolation

Both the cases record responses to defects in audio quality by trials with different frequencies and/or power. They read as follows:

Case (a): This dealt with the problem of planning a broadcast from Karachi to London. The transmission on the usual frequency has been disturbed by sudden high solar activity leading to a geomagnetic storm.

Such occasional high solar radiations, in an otherwise normal situation (as described by the *sunspot number*, see Section 2), unsettle the ionised layers, causing disturbances in transmission. Since the higher layers are more likely to receive the bulk of the additional solar discharge, they are more likely to be unsettled. Hence, an expert hopes that lower-altitude ionised layers may have a better chance of behaving predictably, and therefore decides to see whether an adequate frequency can be allocated in a lower-frequency band.

The attempt recorded in this case was for 6245 kHz. However, it was found that this frequency has been blocked by deliberate jamming. Hence a second attempt was made in the next higher band. In this trial 7025 kHz was chosen. This frequency and the ones around it have turned out to be heavily occupied by many amateur radio operators in Western Europe. The third attempt, on frequency 9133 kHz, however, was found 'ok'. But this frequency is normally reserved for special non-broadcast use, such as aeronautical weather reports and communications. Hence an unblocked legitimate broadcasting frequency reasonably close to this one was sought. The eventual solution was 9440 kHz.

The expert estimated the required power to be 240 kW. However, the maximum power of the available transmitter was 200 kW. Considering that this is significantly smaller than the suggested requirement, boosted-power transmission has been recommended, so another Pakistani transmitter of about 40 kW should be used simultaneously to broadcast the same programme on the same frequency.

Case (d): This referred to a broadcast from Kazakhstan to Frankfurt, in circumstances similar to Case (a), except that the degree of sudden solar activity was described as 'very high'. The broadcaster first determined the appropriate power to be 300 kW. First, 5260 kHz, which is one of the standard frequencies for Kazakh domestic broadcasting, was tried. There was no reception at all. An attempt on a higher frequency, 7470 kHz, was found to give unclear reception. Transmissions at a still higher frequency, 9150 kHz (in the 9 MHz band), were very noisy, possibly due to unusual solar activity. It was therefore understood that good audio quality could not be ensured by direct transmission. Hence a relay was tried. Ankara was selected as the relay centre, as it was between the source and target and the transmission on a 9 MHz frequency was clearly audible there. The consequent actions therefore were to determine the frequencies and powers for each phase of the transmission. The solution obtained was: transmission on 9150 kHz from Kazakhstan to Ankara using power 250 kW; and relay transmission from Ankara to Frankfurt on 11710 kHz using 120 kW. 11710 kHz has been chosen because it is reserved for regular Turkish international broadcasting.

Case interpolation now aims at finding the set of actions (using the solutions of the two selected cases) that will solve the current problem. Since Case (a) has higher weight than Case (d), the case interpolation scheme follows the sequence of actions carried out in Case (a). For each task, it considers those actions, and selects the one that could be carried out successfully. It then selects a corresponding action from Case (d), and uses it to adjust the parameters of the selected action so that it is best suited for the given problem. However, if there is no matching action in the second case (Case (d), here) then the decision needs to be taken on the basis of the contents in the first case alone. Choice must then be made between two modes: the *optimistic* mode, where the method assumes that the solution provided in the first case can be used in the present problem even without any change in its parameters, and the *pessimistic* mode, where the method assumes that the action may not yield satisfactory result for the present problem, and therefore suggests a more cautious action.

The effectiveness of 'case interpolation' is now illustrated by considering the same problem for both 'broadcast' and 'pressing' situations. Details of the algorithm, and other issues, such as case representation are given in (Chatterjee and Campbell, 1997).

5.4.4. Solution for broadcasting

For the problem under consideration (Example 1) the following answer is obtained when the problem is considered to be of 'broadcast' type.

The method first considers 'determining frequency', and finds that a low-end frequency in band 9 should be the best choice. It now tries to select an interpolatable action from Case (d). However, this attempt at selecting a frequency ends in failure. Consequently, the method is forced to make a binary interpolation between the two options—simply whether or not to use band 9. Since the problem refers to a 'broadcast', the method switches into a pessimistic mode, i.e., one that tries to ensure good-quality transmission. The interpolation then suggests that one should not expect a direct broadcast on a single frequency.

The subsequent actions of Case (a) therefore become irrelevant, and the method now makes decisions from Case (d) alone. Hence the next action that it follows up is *use-relay-transmitter* (Section 4.3). A discrete selection method is applied on the list of possible relay centres associated with Calcutta, as obtained from the *linkage-availability table* (Section 4.4). The method orders the possible candidates on the basis of their distances from Ankara, the relay-centre for the Case (d). The most appropriate relay centre is thus identified to be 'Tel Aviv'.

The method now determines the other necessary parameters from Case (d). For transmitting from Calcutta to Tel Aviv, the power is recommended as 240 kW. The frequency for this phase of transmission is recommended as a low-end frequency on band 9. A search of the stored information about scheduled broadcasts suggests the frequency 9425 kHz for this purpose. The ideal relay-power is taken directly from Case (d), which gives 120 kW. The frequency for this phase of transmission is determined to be in the lower part of band 11, and a search for an unblocked frequency results in selection of 11605 kHz. The final interpolated solution therefore reads:

Transmission from Calcutta to London
Time of transmission (UTC): 1115 to 1145
Seasonal-effect: summer
No direct transmission is possible due to
solar activity: high
use-relay-transmitter—relay-centre is Tel Aviv
Transmission from Calcutta to Tel Aviv
Use Frequency 9425 kHz
Use Power 240 kW
Use Transmitter local-transmitter
Transmission from Tel Aviv to London
Use Frequency 11605 kHz
Use Power 120 kW
Use Transmitter local-transmitter

Evidently, the solution can be further improved by referring to appropriate rules/cases for each phase of transmission. However, in the current approach, the level of complication above is enough to demonstrate the essential operations with cache and interpolation.

5.4.5. Case interpolation when 'pressing'

The method proceeds in a similar way to that given for the 'broadcast' situation. However, since the situation is pressing, it operates in an optimistic mode. Hence solutions that would not be admissible for broadcasting are tolerated. For example, not-too-clear audio, transmission with limited power, and excursion to frequencies outside the normal range of allowed frequencies may comprise a solution for 'pressing' situations.

In this example, the method does not recommend a relay; instead it proposes a direct transmission on a low-end frequency in band 9. The recommended frequency is then decided (after searching the *blocked-frequency table*, Section 4.4) as 9425 kHz. The method now tries to estimate the power. Case (a) suggested 240 kW when the solar activity was 'high', while Case (d) recommends 300 kW when the degree of solar activity is described as 'very high', to make a direct transmission. Hence for the current situation (where the degree of solar activity is 'rather high') the method suggests the ideal power to be 190 kW, by interpolating—actually extrapolating, here—on (*very high, 300*) and (*high, 240*) for the value *rather high*.

However, here again it is found that the source does not have sufficient power to match this requirement, as the maximum available power is recorded to be 150 kW. Considering that it is a 'pressing'-type problem, the method suggests simply the use of the maximum possible power, with the warning that the quality of reception is likely to be poor.

5.5. Time-criticality and the cache scheme

Knowledge-based computations are primarily search-dependent and therefore open-ended in nature. Consequently, their application in time-critical computations, where results are expected within a specific time limit, has been somewhat limited.

The caching approach described in this article is intended as a way of answering the demands of time-criticality. It permits a solution or method tailored to the characteristics of a given time-critical request to be found quickly, and even retrieved on demand if there is only enough time for retrieval of a single default solution.

Section 4.4 indicates three criteria by which the effectiveness of a cache can be assessed: successful retrieval of knowledge, maintaining that success in time-critical situations, and progressive improvement in quality of the results as deeper levels of the cache are accessed and the computing time increases. Examples of progressive quality of knowledge-based outputs as deeper levels of the cache are accessed have already been demonstrated. The treatment of effectiveness in the present section completes the assessment by considering time-criticality.

The relevant experiments have consisted of two steps:

- firstly, to estimate the temporal requirements for computations at each cache level.
- secondly, to use those estimations to determine the temporal bounds attached to each cache level, and examine to what extent these bounds are respected in subsequent tests.

The temporal requirements have been measured in terms of overall run time, which includes not only the computing time for the given task, but also the many behind-the-scene operations of a multi-user operating system that are outside a single user's control. In the tests, two kinds of rule-level computations have been distinguished: *Rule1*—when a default plan is carried out; and *Rule2*—when transmission conditions do not allow use of a default plan of direct transmission, so that relaying is necessary. Thus for these experiments the radio cache had five levels. The average computational time observed by repeating computations of each of the five types (*Default1*, *Default2*, *Rule1*, *Rule2* and *Case*) are 312, 600, 645, 1139 and 1553 milliseconds (ms) respectively, with respective standard deviations (SD) being 41.6, 42.7, 47.1, 59.9 and 60.4.

To ascertain the reliability of the cache, temporal bounds for each level have been calculated in the following way: for the *i*th level, compute a lower threshold LT_i as *Average—SD*; and an upper threshold UT_i as *Average+SD*. The greater of UT_i and LT_{i+1} is then considered as an upper bound for level *i*. The upper bound for the $(i-1)$th level will then act as the lower bound for the *i*-th level. The method refuses to provide any answer if the allotted time T is less than the observed minimum time for the shallowest or 1st level. For the deepest level, the observed maximum requirement plus the corresponding standard deviation is taken as the upper limit, although in real use there may not be any upper bound for the deepest level of the cache.

This type of bound selection may be called a *soft* (Dean and Wellman, 1991) approach, since computation with a 'soft' bound tries to ensure that the best possible reasoning technique that is achievable within the given temporal bound is used in solving a given problem. For each of the five types of computations, experiments have been repeated 100 times, with the allowed time T selected randomly for each experiment. The success rates for the five levels have been 80, 42, 92, 67 and 91%, respectively. These rates are not too satisfactory as the definition of an upper bound is too optimistic.

In a second set of experiments, the lower bounds have been chosen more pessimistically using the formula $LT_i = Average + SD*3$, so that the bound should statistically contain 99.7% of the observations. The lower bound of a level *i* was still regarded as the upper bound of the level $i-1$. The success rates under this scheme for the five levels have been 80, 82, 100, 98 and 93%. These are rates that are adequate for some types of soft time-critical requirements. Where an application demands better performance, this can be achieved by readjusting the upper limits associated with the levels.

In (Chatterjee and Campbell, 1998a) details of such experiments and their interpretations have been described more fully. The conclusion drawn from there is that the cache-based knowledge-representation scheme provides a solid foundation for using knowledge-based techniques in time-critical situations. Although the success rate is not 100% in all the situations, the cause can be attributed not to some inherent drawback of the scheme, but to uncertain behaviour in the computing environment. Thus the cache-based scheme may not be suitable for the *hard* (Dean and Wellman, 1991) type of time-criticality, where a failure to abide by the imposed time limits may lead to disastrous effects, but a caching scheme with suitably chosen time bounds for each level appears to be an efficient approach for dealing with time-critical problems of a *soft* nature.

6. General features of a cache-based approach

If one is adopting a cache-based scheme for representing and using knowledge, the normal principles of knowledge acquisition and related areas of artificial intelligence apply, i.e., caching leads to no short cuts in that process. There are also some considerations, dictated by the structure of the cache representation, which may change the emphasis in knowledge acquisition, or call for additional information from the sources that have supplied the existing knowledge.

These considerations are as follows:

- **Default-level actions**, which can be obtained immediately or almost immediately on demand, are a standard part of a typical time-critical system. Experts in a subject where time criticality may occur should be asked explicitly for characterisation of the different types of problem situations where instant responses were feasible, and also, their knowledge of what to do when they are responding instantly to these problems. It can be expected that this knowledge will take the form of *simple algorithms* or the *contents of tables* where suitably-indexed default answers will be stored. If anything different is proffered, the expert may be asked to re-cast the knowledge into one of those two forms so that it can be accommodated in the default levels of the cache.

- **Choice of the number of levels** for a cache involves identification of the number of distinct knowledge-based means, with different expected computing times, of resolving problems in the domain under consideration. Broadly speaking, one should look first for three levels: default, heuristic and 'most exhaustive'. The last one may mean some case-like solution, or some substantial algorithmic programs, as in process control, or approaches that rely on methods of planning from artificial intelligence. Sometimes there may be as few as two levels, if the most exhaustive treatment is via heuristics expressed in rules. On other occasions, as in the radio example, one level of knowledge may split into multiple levels (here, Default1 and Default2) on a closer inspection if it is to take the fullest advantage of the notion of caching. In acquiring and checking acquired knowledge for building a cache in an arbitrary domain, the question 'does this information permit more levels to exist?' is a good focus for attention. The two levels described in Sections 5.1 and 5.2 are practical examples of that question in action.
- **Identification of key features or abstractions** which stake out the space of possible problems is the most important step in preparing a cache to accommodate knowledge. Each such feature may lead to one dimension in the cache, though one should try to be economical and keep the dimensionality as low as possible. These points are best understood through an extended example, which is why the present paper treats the issue at length for the shortwave radio domain. In the early stages of knowledge acquisition for designing a cache, emphasis should therefore be laid upon collecting the key features for problem characterisation.
- **Determination of characteristic values** for each key feature that labels one dimension of a cache follows from the identification of the features themselves. Each of its characteristic values (i.e., V values), which will label one column in the relevant column dimension of the cache, also serves to distinguish one set of situations from all the others. Section 3 explains the background. The sources of expertise should be consulted, as early as possible during knowledge acquisition, to establish this information. Once obtained, it leads to further useful questioning because it determines the structure of the cache. For each cell, one should ask 'What item(s) of knowledge belong here and make this cell different from any other one?' At the case level, in particular, this is effective in eliciting examples that may otherwise be overlooked.
- Since **case-based knowledge** is represented more informally than rules or algorithms, methodologies for knowledge acquisition give it limited attention. Even in exhaustive texts on case-based reasoning (Kolodner, 1993; Aamodt and Plaza, 1994; Voss, 1996), this remains true. Case elicitation is usually treated as a mixture of *collecting anecdotes* and the *seeking of examples* to illustrate when and where more formal knowledge (e.g., rules) breaks down or is insufficient. However, the notion of caching imposes discipline on such enquiries. This is because case acquisition here is guided by the cache structure. The indexing supplied through the key features, their characteristic values, and the question of what to put in each cell of the cache at a case level help in directed and efficacious case elicitation. Furthermore, as the output of a time-critical computation is always read as a plan (an instruction, or sequence of instructions) for resolving a problem, the cache compels a gross form for any case. In this form there are four parts: a description of a problem situation, with the indexing labels and their values represented in a way that allows easy access; a (plan of a) response to resolve the problem; an outcome; all other pertinent information, e.g. on reasons for any unsatisfactory aspects of the outcome. This form is quite a precise guide to eliciting and representing knowledge as cases, while still allowing some freedom over how to fill in their details.

Determining the **characteristic computing time** (T_is, Section 3.1) associated with each level of a cache is also an essential part of the approach. This is the time that a computation at a given level is likely to take. These times are unrelated to knowledge acquisition at the initial stage. When the other parts of the cache are filled, each T_i may be determined initially through experiments. They may later be updated during practical use of the cache.

7. Concluding remarks

The work that is described here shows that knowledge-based computing can be used effectively in solving problems in a decidedly non-trivial application. They satisfy the criteria stated in Section 3.

The novelty of the work should be assessed from the general view point of time-critical applications of knowledge-based techniques. Although the earliest reported application of such techniques can be traced back to 1983 (Masui et al., 1983) the reported state of the art has not advanced much further since then. It is likely that this is due to the *open-ended* nature of knowledge-based computations. Whether their methods are rule-based or case-based, their *modus operandi* can be summarised as searching forward and/or backtracking through a knowledge base until an acceptable solution to the given problem is found. As a result, the overall solving procedure can make unpredictable demands on computing time, and is therefore not well suited to time-critical computations. The cache-based knowledge-representation scheme obviates this open-ended searching, as the process of retrieval from the cache has a more predictably algorithmic character and is therefore open to reliable temporal estimations. That an exactly similar matching may not be retrieved from the cache may seem to be an initial drawback of the scheme. But interpolation responds to this situation, in a way that is well understood from numerical analysis, and which has some psychological validity (Simon, 1959) for human attempts, under pressure of time, to solve problems closely related to examples for which solutions are already known.

Although 'knowledge interpolation' has arisen from the need to make the caching scheme work effectively, it is useful independently of any cache. For example, if a situation S does not match exactly with the preconditions (for forward chaining) or postconditions (for backward chaining) of any rule in a rule base, but it is known from some semantic information that S belongs in the same 'dimension' as terms in at least two of the rules, interpolation can be performed on the two rules closest to S, as described in this article and in (Chatterjee and Campbell, 1998b).

The caching scheme can be treated, further, as a guide to the elicitation or checking of knowledge, particularly case-based knowledge (where existing knowledge-elicitation methodologies are least specific). The labelling of cache columns, which leads to the assignments of labels to all cells, then raises the question 'What knowledge is appropriate to solve a problem described by such a label?' for each cell, and ensures that the elicitation or checking is systematic.

Since, in time-critical problem solving, one's requirement is simply any workable solution that can maintain the temporal bounds, and not the qualitatively best solution (Kopeikina et al., 1988), the cache-based scheme in conjunction with knowledge interpolation provides a practicable answer to this requirement. The extensive treatment of one domain of application, in this article, is intended as a case-study. As none of the features or structures in the method is specific to shortwave radio, it illustrates how to use the notion of caching for time-critical problems in general.

Referenc

Aamodt, A., Plaza, E., 1994. Case-based reasoning: foundational issues, methodological variations and system approaches. AI Communications 7 (1), 39–52.

Bird, M., 1998. Space weather 1998. In: World Radio TV Handbook. BillBoard Publications Inc, New York, pp. 52–57.

van Benthem, H.I., 1995. GRAPH—Generating Radio Link Assignment Problems Heuristically. Master's Thesis, Faculty of Technical Mathematics and Informatics, Delft University of Technology, Delft, Netherlands.

Braun, G., 1982. Planning and Engineering of Shortwave Links. John Wiley and Sons, Siemens Aktiengesellschaft, Chichester.

Chatterjee, N., Campbell, J.A., 1994. Adaptation through interpolation for time-Critical case-based reasoning. Topics in Case-Based Reasoning. In: Wess, S., Althoff, K., Richter, M.M. (Eds.), Lecture Notes in Artificial Intelligence, 837. Springer-Verlag, Berlin, pp. 221–233.

Chatterjee, N., Campbell, J.A., 1996. Knowledge interpolation: a new approach to rapid symbolic reasoning. In: Anjaneyulu, K.S.R., Sasikumar, M., Ramani, S. (Eds.), Proc. Knowledge-Based Computer Systems: Research and Application. Narosa Publishing House, pp. 67–78.

Chatterjee, N., Campbell, J.A., 1997. Interpolation as a means of fast adaptation in case-based problem solving. In: Bergmann, R., Wilke, W. (Eds.), Proc. 5th German Workshop on Case-Based Reasoning. Centre for Learning Systems and Applications, University of Kaiserslautern, Germany, pp. 65–74.

Chatterjee, N., Campbell, J.A., 1998a. Cashing in on caching: an architecture for time-bounded knowledge-based problem solving. Real-Time Systems 15 (3), 221–247.

Chatterjee, N., Campbell, J.A., 1998b. Knowledge interpolation: a simple approach to rapid symbolic reasoning. Computers and Artificial Intelligence 17 (6), 517–551.

Dean, T.L., Wellman, M.P., 1991. Planning and Control. Morgan Kaufmann, CA.

Jacobs, G., 1992. Do-it-yourself propagation forecasting table for the most suitable Megahertz broadcasting bands. In: World Radio TV Handbook, 46. BillBoard Publications Inc, New York, pp. 30–32.

Kolodner, J.L., 1993. Case-Based Reasoning. Morgan-Kaufmann, CA.

Kopeikina, L., Brandau, R., Lemmon, A., 1988. Case-based reasoning for continuous control. In: Proc. DARPA Case-Based Reasoning Workshop. Morgan Kaufmann, CA, pp. 233–249.

Magne, P.E., 1998. Passport to World Band Radio. IBS Ltd, Penns Park, Pennsylvania.

Masui, S., McDermott, J., Sobel, A., 1983. Decision making in time-critical situations. In: Proc. IJCAI-83. William Kaufmann Inc, Los Altos, CA, pp. 233–235.

Simon, H.A., 1959. Theories in Decision making in Economics and Behavioral science. American Economic Review 49, 253–283.

Sneath, P.H.A., Sokal, R.R., 1973. Numerical Taxonomy: The Principles and Practice of Numerical Classification. W. H. Freeman and Co, San Francisco, CA.

Voss, A., 1996. How to solve complex problems with cases. Engineering Applications of Artificial Intelligence 9 (4), 377–384.

Zadeh, L.A., 1979. A theory in approximate reasoning. In: Hayes, J.E., Michie, D., Mikulich, L.I. (Eds.), Machine Intelligence, 9. Halstead Press, New York, pp. 149–194.