APPLICATION OF KNOWLEDGE BASED SYSTEMS TECHNOLOGY TO TRIPLE QUADRUPOLE MASS SPECTROMETRY (TQMS)

Hal R. Brand and Carla M. Wong

Lawrence Livermore National Laboratory
P.O. Box 808, Mail Stop L-365, Livermore, Ca. 94550

ABSTRACT

The complexity of chemical instrumentation is such that automation of certain instrument functions by conventional algorithmic means is either very difficult or completely unsuitable. This paper details work in progress on the application of knowledge based systems technology to the tuning of a complex analytical instrument, a triple quadrupole mass spectrometer (TQMS). The knowledge representation schemes and interface design between the expert system and the TQMS instrument are discussed. Preliminary results of optimizing the TQMS on chemical standards are presented.

1. Introduction

In the past twenty years, chemical instrumentation has become very powerful and complex. This complexity occurs not only in the operational principles and physical construction of the instrument, but also in the acquisition and interpretation of data. In many cases, micro- or mini-computers are required for the operation of the instrument; and in most of these, a dedicated computer collects data, transforms the data when necessary and displays the results.

This automation has traditionally been achieved using standard algorithmic methods implemented in procedural languages such as FORTRAN, FORTH and assembly code. However, these techniques are not suitable when the automation task involves poorly specified heuristics such as tuning or optimizing the operational parameters for very complex instruments or processes. Recent availability of commercial knowledge based systems tools has made the automation of these problems possible. This paper presents work in progress on the application of knowledge based systems techniques to the tuning of an analytical chemistry instrument, a triple quadrupole mass spectrometer (TQMS).

"Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under contract number W-7405-ENG-48."

2. The TQMS Domain

A triple quadrupole mass spectrometer is a sophisticated chemical measurement instrument that has been described in detail elsewhere (Yost and Enke, 1978; Yost, et al, 1979; Yost and Enke, 1979). It can be thought of as a two stage mass spectrometer with a collision cell between the two stages (figure 1). As in a normal, single stage mass spectrometer, a sample is introduced into the source where it is ionized and the fragments are accelerated into the first quadrupole. This quadrupole can be set to act as a mass filter, allowing only those ions with a specific mass to charge ratio to pass into the second quadrupole. If the instrument is operated in normal mass spectrometry (MS) mode, then all of the ions are detected and a mass spectrum (a plot of intensity vs. mass to charge ratio) is produced (figure 2). If the second quadrupole region is used as a strongly focusing reaction chamber and is pressurized with an inert gas, the "parent" ions selected by the first quadrupole collide with the inert gas molecules and are further fragmented to form "daughter" ions. These daughter ions are then mass selected in the third quadrupole and analyzed as in normal mass spectrometry producing a mass spectrum of a mass spectrum (MS/MS mode).

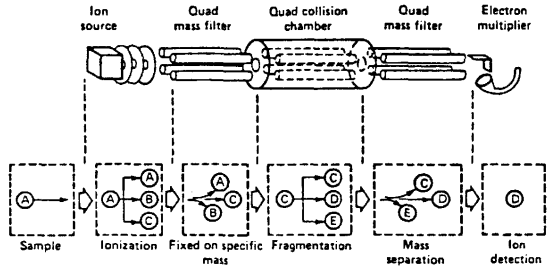

Operation Mode	Quad 1	Quad 2	Quad 3	Results
1	Separated by mass	All masses passed No gas	All masses passed	Normal mass spectrum
2	Fixed on specific mass	All masses passed Collision gas	Separated by mass	Spectrum of all daughter ions from the selected parent ion
3	Separated by mass	All masses passed Collision gas	Fixed on specific mass	Spectrum of parent ions that fragment to give specific daughter ion
4	Separated by mass	All masses passed Collision gas	Separated by mass	Fixed mass difference between 2 scanning quads gives specific neutral mass loss
5	Fixed on specific mass	All masses passed Collision gas	Fixed on specific mass	Single or multiple reaction monitoring

Figure 1. TQMS Schematic

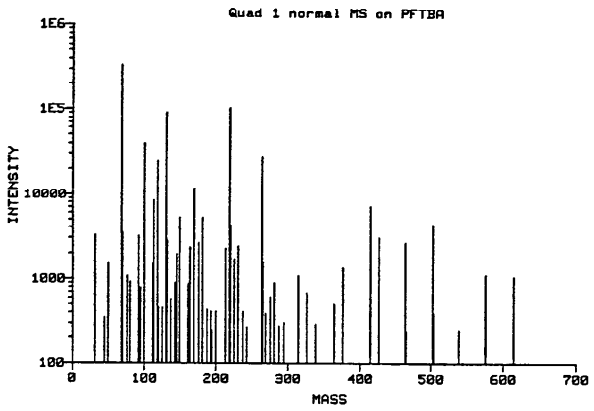

Figure 2. Normal Mass Spectrum (MS mode)

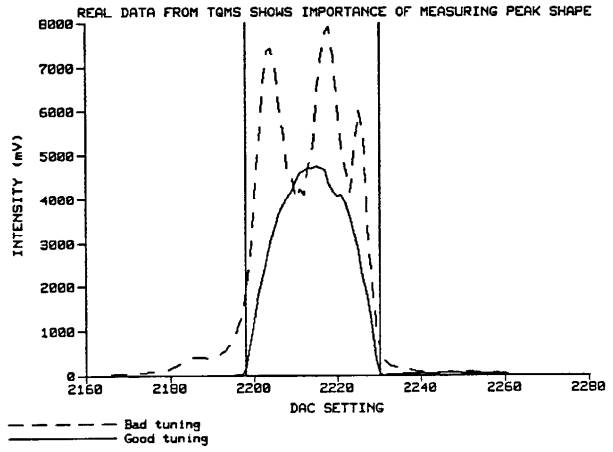

Figure 3. Good and Bad Peak Shape Comparison

Each stage of this process (ionization, mass filtering, collision in the second quadrupole and another stage of mass filtering) is inefficient. In order to obtain maximum sensitivity and selectivity from this instrument, each of these inefficiencies must be minimized by careful tuning of the instrumental parameters in either MS or MS/MS mode.

The TQMS built at Lawrence Livermore National Laboratory has over 30 operational parameters controlled by a DEC LSI-11/23 micro-computer (Wong et al, 1983). The LSI-11 is programmed to acquire data for the chemist after the instrumental parameters have been manually optimized (tuned). The tuning of the TQMS is a labor intensive process, requiring at least 30 minutes to obtain a parameter set where the resulting "tune" is a compromise over the mass range of interest. Studies have shown that the ion intensities may be increased by factors of 2 to 30 if the instrument is tuned over many small mass ranges instead of one large range. For routine analyses, however, the operator time required to tune for many mass ranges cannot be justified.

The process of tuning the TQMS is simply an optimization problem with many independent parameters. The operational parameters of the instrument are varied until optimum sensitivity is obtained within the constraints that peak shape must be "good" and adjacent peaks must be resolved. This translates to maximizing peak height while maintaining a nearly parabolic peak shape with a peak width of less than one mass unit (Wong, Kunz, and Kehler, 1984). Figure 3 demonstrates the effect of one parameter on peak shape, and shows both an acceptable and unacceptable peak shape.

The manual process of tuning the TQMS consists of the expert adjusting the operational parameters of the instrument while watching an oscilloscope display of several peaks distributed over the mass range of interest. The operator then maximizes peak heights while ensuring that peak width and shape remain within the constraints. This process creates an average tune over the entire mass range where the sensitivity of each ion is compromised (Wong, Kunz, and Kehler, 1984; Wong, Crawford, et al, 1984; Wong, Lanning, et al, 1984).

The ability to automate the tuning procedure would allow the chemist to tune the TQMS over small mass ranges (or even tune for specific parent/daughter combinations) to obtain increased sensitivity for every ion. To accomplish this, we tried two approaches: an algorithmic approach and a expert system approach. Experimentation indicated that the time required to tune the instrument using the algorithmic approach (in this case, a Simplex algorithm) was excessive due to the data acquisition time. A manual tuning required approximately 30 minutes while the Simplex tuning required approximately 300 minutes (figure 4). Comparing the Simplex approach with manual tuning revealed that the Simplex algorithm could not take advantage of any knowledge of the TQMS that would permit acquiring less data. For example, each step of the Simplex algorithm adjusts every instrumental parameter which

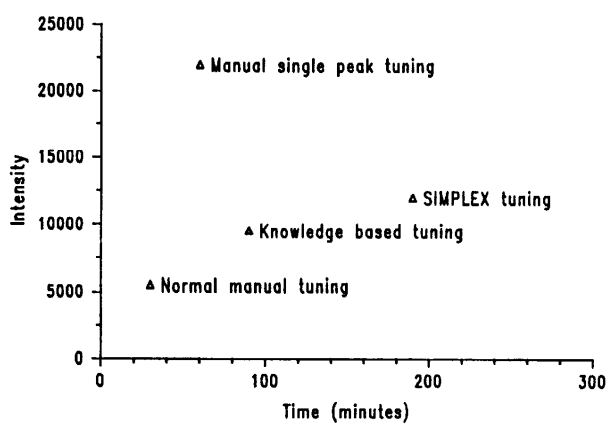

Figure 4. Comparison of Tuning Methods

requires many data points to be taken in order to evaluate the effect on peak shape and therefore on overall performance. The human expert is able to chose which knob should be adjusted for the maximum desired effect and knows whether adjusting that knob will affect the peak shape or merely the peak height. Since many of the knobs do not affect peak shape under the right circumstances, adding this knowledge into the system leads to acquiring less data (single point peak height data instead of many point peak shape data), and the entire tuning process proceeds more rapidly (Wong and Brand, 1985).

3. Knowledge Representation

Three major classes of knowledge are relevant to tuning the TQMS:

1. Knowledge about the instrument itself.
2. Knowledge about how to evaluate the instrument output and tune the instrument.
3. Knowledge about how to interface to the LSI-11 computer to control the instrument.

The knowledge engineering tool KEE*, running on a Xerox 1109 LISP processor, was used to develop this expert system. KEE provides four basic mechanisms for representing and using knowledge: frames, methods, rules, and active values. Frames permit static knowledge to be represented in a class/sub-class/member inheritance hierarchy. Methods provide a mechanism for representing procedural knowledge in LISP code with an object-oriented interface and control structure. Rules provide an alternative mechanism for representing procedural knowledge with KEE providing both forward chaining and backward chaining control structures. Finally, active values provide a mechanism for attaching procedural knowledge in the form of methods or rules to frames.

The hybrid knowledge representation environment provided by KEE permits significant flexibility in the representation of knowledge about the instrument and its tuning procedure. Frames provide a very clear and simple representation of the static (or declarative) knowledge about the instrument parts and controls (see figures 5 and 6). Knowledge about the attributes of the objects or classes represented by the frames are represented in slots within the frames. An inheritance mechanism allows specification by differences, greatly simplifying the job of specifying and maintaining the knowledge in the frames. The class/sub-class/member relationships of the frames provide another mechanism for encoding knowledge with frames.

Methods are valuable for representing procedural knowledge about the time sequencing of the steps in tuning the TQMS, and for implementing the standard algorithms necessary to interface the expert system to the LSI-11 control computer. A flexible interface between methods and the rule system allows the knowledge of the

*Trademark of IntelliCorp

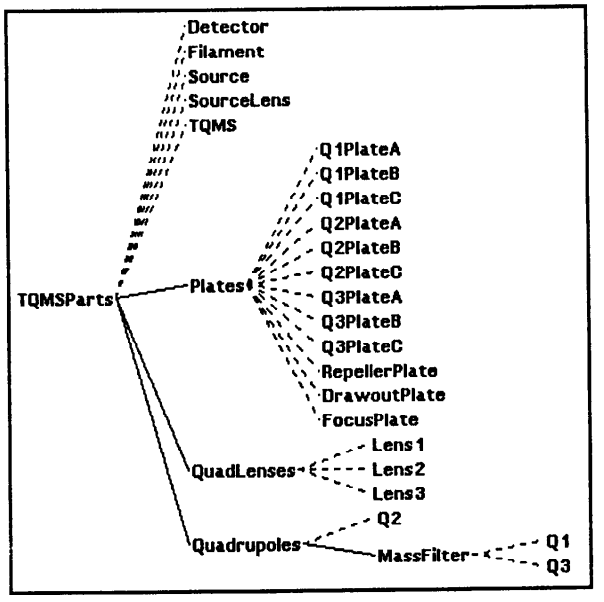

Figure 5. TQMSParts Hierarchy

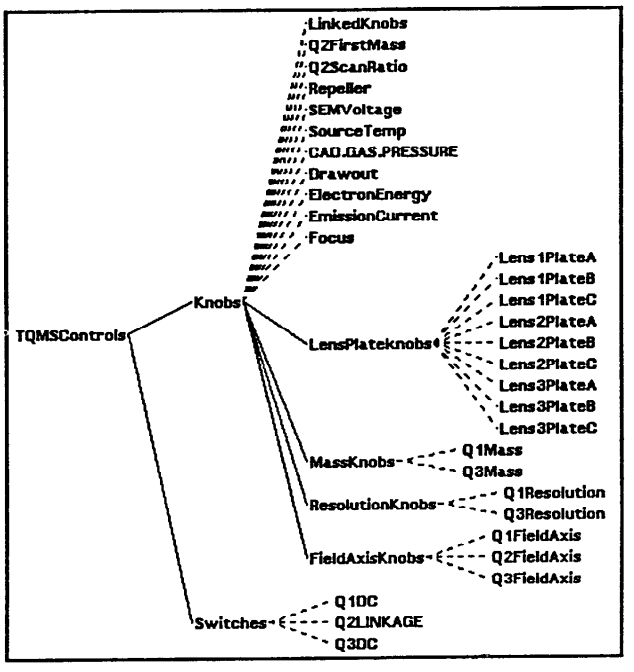

Figure 6. TQMSControls Hierarchy

iterative steps in tuning the TQMS to be simply and clearly represented in a method, while concurrently permitting the use of rules for the complex decision steps.

Rules, with a backward chaining control structure, are used to represent the knowledge of which knob to adjust or "tweak". This largely heuristic and poorly defined knowledge was significantly easier to represent in rules for

two reasons. First, rules provided a procedural knowledge representation scheme that was readily understood by the experts which facilitated information transfer from the experts to the knowledge based system. Second, using rules provided for the incremental addition of knowledge about knob selection since new rules could be added without regard to their placement or order of use. Rules allowed the experts to concentrate on expressing knowledge about the TQMS and not on the decision control structure.

A natural representation for the value of an instrument control parameter, or knob, is a "Setting" slot in the frame that represents the knob. Active values permitted this representation scheme by providing a mechanism that associates the procedural knowledge about interfacing to the LSI-11 computer with the "Setting" slot of each knob frame. By placing an active value on the "Setting" slot of each knob frame, methods are invoked at each access to the "Setting" slot. These methods cause the instrument's physical knob settings to track the "Setting" slot in the knob frames. The advantage of this interfacing scheme is the invisibility of the instrument interface to the rest of the system thereby permitting continued development and testing of the system during times when the instrument is not available for development work.

3.1 Representation of the Instrument

Knowledge about the physical construction of the TQMS in terms of parts and assemblies is represented with frames as shown in figure 5. Frames are also used to represent the physical controls (knobs and switches) on the TQMS (figure 6). Slots within each hierarchy are used to represent the knowledge about which knob(s) and/or switch(es) control which part(s)/assemblies. In addition, slots within the TQMSParts hierarchy represent the part/assembly relationships. This representation scheme clearly separates the instrument knobs and switches from the instrument parts and assemblies. The first attempt at representing the instrument with frames did not make this distinction as the experts very often blur the two together because there is often a one-to-one correspondence between the part being controlled and the control(s) of that part. The source of this confusion can be understood by examining figure 7, which was generated by methods within the TQMSParts frames that interpret the part/assembly and part/control links. In this figure bold faced print indicates the knob(s) that control the parts/assemblies shown in normal print. In many cases there is a one-to-one correspondence between TQMSParts and Knobs (i.e. the RepellerPlate part and the Repeller knob), but there are also cases where there is a one-to-many relationship (i.e. the Q3 part controlled by the knobs Q3Mass, Q3Resolution, and Q3FieldAxis). This caused some frames to represent a part/knob pair, while similar frames represented only a single knob. This led to difficulties with the active value interfacing scheme and inheritance. Separation of this

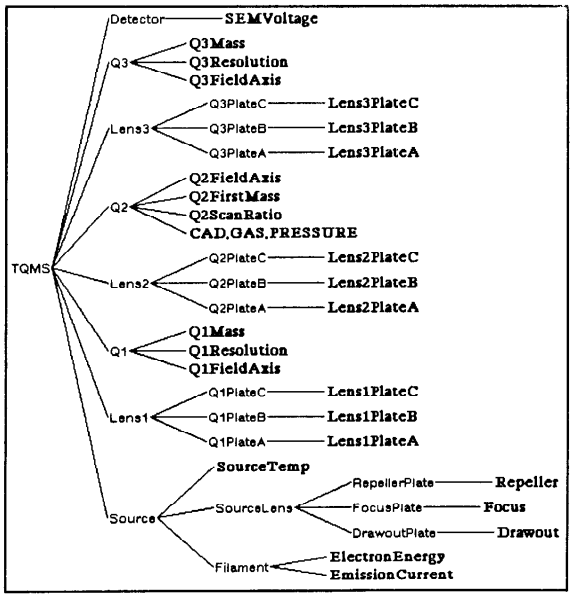

Figure 7. TQMS Part/Whole Graph

knowledge into two inheritance hierarchies with accompanying methods for producing the TQMS part/whole graph improved the transparency of the representation, increased the usefulness of inheritance, and provides a clear depiction of the part/whole breakdown of the instrument which is the way experts view the TQMS.

The chosen representation also makes possible the use of "virtual knobs" (see "LinkedKnobs" in figure 6). A "virtual knob" may be defined to control two or more TQMS knobs (and therefore one or more parts) simultaneously, giving them a single setting as though they are controlled by a single physical knob. The mapping from the setting of the virtual knob to the settings of the physical knobs is dependent upon the type of virtual knob implemented. The TQMS experts had determined that the settings of certain physical knobs should be varied together in a fixed way, but they were unable to do this effectively while manually tuning the instrument. Virtual knobs provide a simple and effective mechanism to accomplish this task.

3.2 Representation of the Tuning and Evaluation Procedures

Knowledge about the process of tuning the TQMS is represented using methods and rules. A method is used to represent the high level procedure of tuning the TQMS which iterates over the following steps:

1. Use the current output of the instrument and the history of what has already been done to select a knob to adjust; if a knob cannot be selected, the tuning process is complete.
2. Determine the parameters necessary to "tweak" (adjust) that knob.

3. Tweak the knob to achieve an increase in instrument performance.

This procedure is represented as a method because it is well specified and not believed to be subject to significant change. While methods are often less transparent to the experts than rules, it was felt that the convoluted rules necessary to represent this iterative, step-wise procedure would only serve to obscure the simplicity of the process. In addition, much of the bookkeeping details underlying this process would have to be implemented in LISP code within the rules, further reducing the clarity of the rules.

Rules are used to represent the heuristic decision-making knowledge needed by the first two steps of the tuning procedure. The clarity of the rule structure made incremental improvement of the knowledge possible. Backward chaining from the hypothesis "THE KNOB.TO.TWEAK OF CURRENT.TUNE IS ?X" was selected as the control structure for applying the rules for knob selection. The same control structure starting from similar hypotheses is used with the rules that determined the tweaking parameters in step two. Backward chaining was chosen because it limited the inferences made to those necessary to make the required decision, and allowed for terminating the chaining process when the decision was made. The forward chaining mechanism in KEE was discarded because it did not provide any mechanism to terminate forward chaining once the required decision was made.

Step two of the tuning procedure requires five separate decisions. These decisions are largely independent and the rules used to make them are separated into rule sets and invoked sequentially following the choice of the knob to tweak in step one. The coupling between these decisions is handled by permitting any rule set to make a decision that would normally be made later by another rule set. When this happens, the rule set corresponding to that decision is skipped since the decision has already been made. For example, if the rule set that selects which knob to tweak also specifies the limits of the knob adjustment, the rule set that normally determines the limits is skipped. The advantage to this sequential decision-making process is that any of the subsequent decisions can be made when unusual circumstances are recognized, and the routine decisions can be deferred to the rule sets that handle the usual cases.

The method implementing the tuning procedure provides for separate, modifiable, default decisions should the rule sets fail. These defaults are stored in the frame that represents knowledge about the progress and state of the tuning process. These default decisions reduce the rules needed within each rule set to only those rules necessary for the "exceptional" cases. These defaults also permit the system to function, albeit not always optimally, under circumstances not previously considered.

Methods are used to represent the knowledge of how to tweak a selected knob. The methods implement algorithms for one dimensional, noise insensitive optimization. Methods were chosen for three reasons: 1) such algorithms already existed and were relatively simple to implement as methods; 2) the heuristics associated with tweaking a knob were simple and could be incorporated into the existing algorithms as parameters; and, 3) should the algorithmic approach to representing knowledge about the tweaking procedure be successful, the algorithms could be easily ported to the LSI-11 control computer resulting in a significant increase in instrument tuning speed. The parameters to the tweaking algorithm are selected (as mentioned above) during the ordered decision making process immediately preceeding the tweaking of a knob.

The final knowledge required to tune the TQMS is how to evaluate the signal from the instrument. In the knob selection process, output from the TQMS needs to be analyzed to determine which knob, when properly adjusted, will most likely produce the greatest increase in instrument performance. During the tweaking step, constant evaluation of the instrument output is required to properly adjust the knob. Supporting the decision of which knob to tweak, rules are used to evaluate a condensed description of the peak height, width, and shape. Rules were chosen because they provide a flexible mechanism for evaluating the different factors in the TQMS output at a time when complete evaluation is critical. The evaluation rules are placed in the rule sets that select the knob to tweak, allowing them to be tailored to their companion knob selection rules.

During the tweaking process, the state of the instrument isn't critical and only a single indicator of instrument performance is required. Accordingly, a method is used to convert the condensed peak description into a single numeric performance measure. Use of a method at this stage was motivated by the desire to eventually port the tweaking procedure to the LSI-11 control computer thereby achieving shorter tuning times.

3.3 Representation of Interfacing Knowledge

The knowledge of how to interface the TQMS expert system to the LSI-11 control computer is partitioned into two pieces. The first part of the interface is represented using frames and methods, and the hierarchy of these frames is shown in figure 8. The frames and methods represent the knowledge needed to command the LSI-11 computer to manipulate any of the TQMS controls and to solicit any acquired/processed data from the LSI-11 computer.

Active values are used to connect the interfacing knowledge with the knowledge about the physical controls of the TQMS. Each member frame shown in figure 8 (connected by a dashed line) contains methods that respond to read and write accesses to a slot (or set of slots) in the member frames of the TQMSControls hierarchy (figure 6) via the active value mechanism. The methods contain the knowledge of how to translate

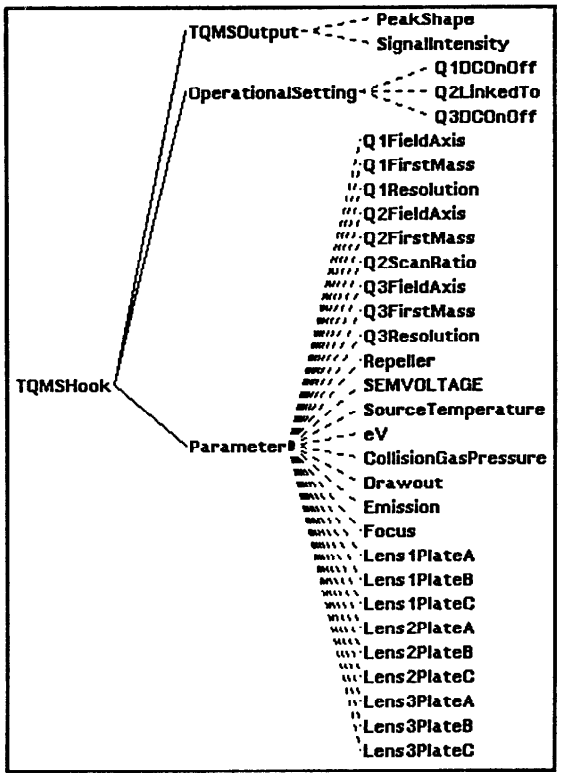

Figure 8. REALTQMS Interface Hierarchy

the read/write request to a slot into the corresponding request/command to the LSI-11 control computer, and how to transmit that request/command. Additional methods are used to connect and disconnect the active values. In the disconnected state, accessing slots of the frames in the TQMSControls hierarchy will not affect the instrument. In the connected state, changing a knob's "Setting" slot results in the instrument being commanded to make the corresponding change, while reading the slot results in the instrument being interrogated for the current value.

The second part of interfacing knowledge consists of LISP code that implements a low-level master/slave protocol to provide an interlocked, reliable communications path. LISP code was chosen because of the need for efficiency dictated by the Xerox 1109's RS-232C interface. The implementation uses hierarchical control and state machine emulation to reliably transfer commands and responses between the Xerox 1109 and the LSI-11 control computers. The RS-232C communication link was chosen because the interface was available on both machines and high bandwidth was not required.

The separation of the interfacing knowledge from the representation of the instrument simplifies experimentation with different instrument representation schemes. The active value separation scheme employed is also ideally suited to the use of an instrument simulator since a simulator can be connected with active values in an identical fashion. Lastly, the ability to disconnect from the instrument facilitates system development and debugging during times when the instrument is not available or unnecessary. The disadvantage of the active value scheme is that significant bookkeeping is required to provide the correspondence between the instrument representation frames (figure 6) and slots and the instrument interface frames (figure 8). The later problem has so far been addressed by a combination of methods and an extra slot in the TQMSControls frames which contains knowledge about connecting to the interface frames of figure 8.

4. Results

The use of knowledge based systems techniques to automate the tuning of the TQMS has proven to be very useful. The knowledge based system interfaces to the instrument and exercises intelligent control over the tuning process in real-time. Initial results have demonstrated that the system is able to tune the instrument in MS mode nearly as well as a Simplex optimization procedure (in one half the time) and better than an expert operator does in twice the time (figure 4). If the human expert optimizes on a single peak, manual tuning takes less time than the expert system and can attain twice the sensitivity. However, experts do not individually tune every peak in a mass region because it takes too much time; so a more valid comparison of the system performance is shown in figure 9. Optimizing the instrument in four separate mass regions (less than 100, 100-200, 200-350, greater than 350) has enabled us to increase the peak intensity (and instrument sensitivity) in all regions by factors of 2 to 30. As more rules are added to the system and the current rules are optimized, the sensitivity should increase most noticeably in the high mass region (above mass 500) (Wong, Crawford, et al, 1984; Wong, Lanning, et al, 1984; Wong and Brand, 1985). Having demonstrated the usefulness of knowledge based systems to MS tuning, we turned to the more general and difficult optimization problem, that of MS/MS tuning.

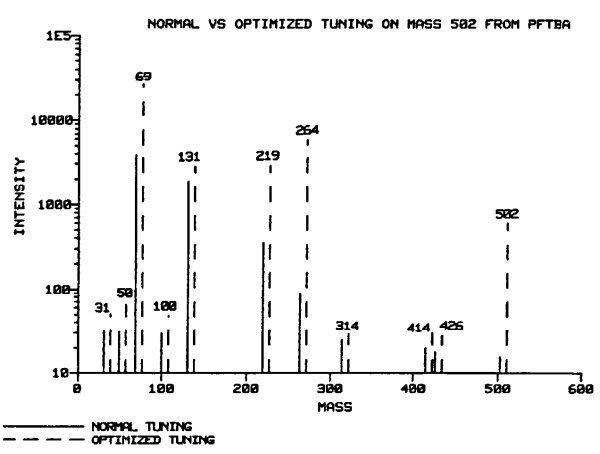

Figure 9. Comparison of Normal and Optimized Tuning

MS/MS operation of the TQMS differs from single MS mode in that selected parent ions are further fragmented and mass analyzed before being detected. This collision process introduces new parameters and conditions which don't exist in single MS operations. For example, the energy of the collision in the second quadrupole is a new parameter to optimize. Figure 10 is a plot of intensity (of the daughter ion at mass 219 from the parent ion at mass 502 from perflurotributylamine, PFTBA) vs. the collision energy. The lower intensity curve shows a typical energy profile which could be obtained by manual optimization of this instrument parameter. By using a rule-based virtual knob to link several of the parameters together, a dramatic increase in the sensitivity was obtained (a factor of 40). This increased sensitivity was never obtained by manual tuning methods, but is easily accomplished with this automated optimization scheme.

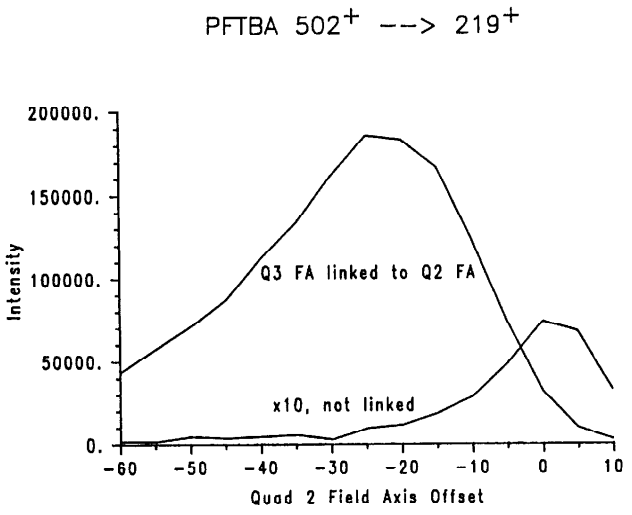

Figure 10. Comparison of Peak Intensity from Linked and Unlinked Tuning Parameters

The current MS/MS optimization system is able to tune the entire instrument to approximately 85% of the sensitivity of a manual tune. As with the single MS tuning, the sensitivity gain of the MS/MS tuning is expected to rise rapidly as rules are added and debugged. Off-loading the tweaking algorithms to the LSI-11 control computer is expected to significantly decrease the tuning time.

An example of the usefulness of the AI technique for tuning the TQMS can be found in a routine analysis of 12 different sulfur compounds in oil shale pyrolysis gas. The expert cannot spend the time it would take to manually tune for each parent/daughter pair. If tuning for each pair took only 10 minutes (an optimistic estimate, as it usually takes about 30 minutes), two hours would be required to optimize the instrument for this analysis. The AI tuning system described in this paper is faster and does a better job of optimizing as it can use virtual knobs unavailable for manual tuning. This frees the operator of the tuning, allowing him/her time for sample preparation, office politics or a cup of coffee.

5. Conclusions

The use of a knowledge engineering tool that integrates multiple knowledge representation schemes significantly decreased the system development time compared to tools that offer only rules. By permitting knowledge to be represented and used in multiple ways, it was possible to select a representation scheme that made the knowledge visible, clear, and easily encoded. The result was enhanced feedback from the experts because they were able to see where the knowledge was stored and how it was encoded.

The interfacing of the expert system to the TQMS proved the value of expertise, encoded in the form of rules, to complex optimization problems. The system is able to optimize the output of a complex instrument running chemical compounds in a way not practical with manual methods. Because the expert system approach allows the instrument to be tuned quickly, multiple mass range, or even individual mass pair tuning is now practical, resulting in large gains in instrument sensitivity.

Separation of the knowledge about knobs and how they are tuned from the knowledge about the TQMS domain makes many aspects of this system transportable to other tuning problems. Much of the knowledge representation techniques applied to the TQMS tuning process could be easily applied to other problem domains. An accelerator, a laser system, and many other complex physical and chemical instruments require that significant time be spent by experts to assure the system is properly tuned to fine working order. Such systems are prime candidates for the application of the same or very similar AI techniques as used on the TQMS to the automation of their tuning process.

Two significant problems with the application of knowledge based systems to chemical instrumentation have been encountered. First, there is a significant learning curve associated with applying the technology, and second, knowledge based systems software tools and the supporting hardware are expensive. While our experience has shown that the hardware and software tools are cost effective for developing the system, these high costs make fielding multiple copies of the system on the development vehicle economically unattractive. The alternative is to port the developed system to other hardware using conventional languages, but this approach is practical only where the large costs of porting the software can be amortized over many systems.

6. ACKNOWLEDGEMENTS

The authors wish to thank Hugh Gregg and Richard Crawford for their valuable programming efforts on the LSI-11 control computer of the TQMS. We also wish to thank Peter Friedland of Stanford, and Hugh Gregg for numerous helpful comments and suggestions on this paper. Special thanks also go to Charles Bender, former Chemistry and Materials Science Department Head at LLNL, who provided the initial funding and encouragement to get this project started, and to Art Lewis for support from the Oil Shale program.

REFERENCES

[1] Yost,R.A., Enke,C.G., "Selected Ion Fragmentation with a Tandem Quadrupole Mass Spectrometer", J. Amer. Chem. Soc., 100, pp. 2274-2275, (1978).

[2] Yost,R.A., Enke,C.G.,McGilvery,D.C., Smith,D., Morrison,J.D., "High Efficiency Collision-Induced Dissociation in an RF-Only Quadrupole", Int. J. of Mass Spectrom. Ion Phys., 30, pp. 127-136, (1979).

[3] Yost,R.A., Enke,C.G., "Triple Quadrupole Mass Spectrometry for Direct Mixture Analysis and Structure Elucidation", Anal. Chem., 51, pp. 1251A-1264A, (1979).

[4] Wong,C.M., Crawford,R.W., Barton,V.C., Brand,H.R., Neufeld,K.W., Bowman,J.E., "Development of a Totally Computer-Controlled Triple Quadrupole Mass Spectrometer System", Rev. Sci. Instrum., 54, pp. 996-1004, (1983).

[5] Wong,C.M., Kunz,J.C.,Kehler,T.P., "Application of Artificial Intelligence to Triple-Quadrupole Mass Spectrometry", IEEE Transactions on Nuclear Science, Vol. NS-31, No. 1, pp. 804-810, (1984).

[6] Wong,C.M., Crawford,R.W., Lanning,S.M., Brand,H.R., "Artificial Intelligence for Optimization of Real-Time Data from Pyrolysis Experiments", 32nd Annual Conference on Mass Spectrometry and Allied Topics, May 27-June 1 1984, San Antonio, Texas, pp. 372-373.

[7] Wong,C.M., Lanning,S.M., Brand,H.R., Crawford,R.W., "Application of AI Programming Techniques to Development an Expert System to Tune a TQMS", 32nd Annual Conference on Mass Spectrometry and Allied Topics, May 27-June 1 1984, San Antonio, Texas, pp. 642-643.

[8] Wong,C.M., Brand,H.R., "Artificial Intelligence: Expert System for Acquisition and Interpretation of Data in Analytical Chemistry", Pittsburgh Conference on Analytical Chemistry and Applied Spectroscopy, Feb. 25-Mar. 1, 1985, New Orleans, La., p. 127.