Qualitative Euler Integration with Continuity

J. Allan Scott

Department of Economics Heriot-Watt University Edinburgh EH14 4AS Scotland UK.

Abstract

The subject of this paper is a novel synchronous fuzzy qualitative simulator developed within the *Mycroft* fuzzy qualitative reasoning framework. Synchronous fuzzy qualitative simulation involves replacing the transition rules of *Mycroft* with an integration phase utilising a qualitative version of Eulers first order approximation to the Taylor series: Qualitative Euler Integration (QEI).

simulation process The described utilises constraint-based fuzzy qualitative models, the variables of which take their values from a predefined fuzzy quantity space. The simulation proceeds, driven by an externally defined integration time step (chosen to ensure the continuity of the magnitudes of the state variables), by means of an explicit Euler integration operation. This provides the set of possible successor values for the magnitudes of the state variables. After this the constraints of the model are solved to provide the values of the non-state variables of the model. As each constraint is solved the same transition rules as for asynchronous simulation are applied to constrain the generation of the behaviour tree. At the end of this process a number of successor states will be generated This number will be less than or equal to the number generated by semi-constructive or non-constructive simulators such as Mycroft or FuSim, and a great deal less than if the transition filters had not been applied. The advantage of this approach is that it permits the utilisation of multiple precision models in which the information concerning the values of system variables may be expressed in vague terms but with precise time stamp information. The system has already been utilised in a research exploring the use of qualitative models for parameter identification, diagnosis, (Steele and Leitch 1997) and control (Keller and Leitch 1994).

Introduction

One of the original motivations for the development of qualitative reasoning systems was a research programme to enable expert systems to reason from first principles, in order to overcome the weaknesses inherent in the first generation, rule-based, expert systems (Kuipers 1986). As the field developed, it became apparent that the programme had problems of its own to overcome if it was to meet its original goals. The dominant problem in qualitative reasoning is that of spurious behaviour generation; and a great deal of research effort has gone into overcoming the problem (Fouche and Kuipers 1992, Kuipers et al. 1991, Lee and Kuipers 1988). In the interim, the engineering community has become interested in, and contributed to the field, because it is seen as a useful tool for simulating the behaviour of complex but incompletely specified plant. Both these influences have contributed to the utilisation of semi-quantitative information (Berleant and Kuipers 1997) and more constructive simulation techniques (Wiegand 1991).

A complement to this, and a major motivation for the work reported herein, is the requirement that a qualitative simulation system be able to match the behaviour of a real system at distinct, measured, timepoints. The reason being that it is possible for data to be presented qualitatively, and yet be precisely time stamped, for example in patient records. For this requirement to be met (with incomplete models) the variables must have, at least, semi-quantitative values (for the magnitudes) in order for temporal intervals to be calculated. In related research the focus has been on the co-operation of symbolic, asynchronous, non-constructive systems with semi-quantitative information, or on the use of constructive synchronous techniques, which have converged on the domain of interval simulation. And with this latter approach, if there is a mapping back into a predefined quantity space, then it is an output operation and the structure of the quantity space is not used to constrain the behaviour generation. In the system presented in this paper, on the other hand, the features of both asynchronous simulation (albeit in the context of fuzzy qualitative simulation (Shen and Leitch 1993, Coghill 1996), and synchronous Euler simulation are utilised to generate behaviours. This is a novel feature of the system presented and is part of an ongoing research program to utilise qualitative simulation for system identification and parameter estimation from time series data, at multiple precisions (Steele and Leitch 1997); and reason with explicit linguistic descriptions

George M. Coghill Department of Computer Science Unviersity of Wales

Aberystwyth SY23 3DB Wales UK. gmc@aber.ac.uk of the data. The approach presented, forms part of the *Mycroft* framework; it has been utilised in a number of contexts already (Keller and Leitch 1994, Steele and Leitch 1997) and the reader is directed to these reports and papers for further information and results.

The structure of this paper is as follows. In the next section the Mycroft framework is summarised, in order to put the subject of this paper in its overall context. In that section the three components required for the qualitative simulation are described: the quantities (in this case fuzzy intervals and numbers), the constraint-based model representation, and each of the other (asynchronous) inference approaches utilised within Mycroft. This is followed by a description of the Euler integration approach and its relation to the use of continuity in the transition analysis and causal propagation phases of the simulation. In section four the technique is further elucidated by comparing the results obtained by utilising the different algorithms from within Mycroft (synchronous and asynchronous). Finally, a selection of similar techniques reported in the research literature are briefly reviewed.

The *Mycroft* Framework

The *Mycroft* framework (Coghill 1996) is a constraintbased fuzzy qualitative reasoning system containing a number of simulation and envisionment algorithms. The development of this framework has permitted the suitability of different techniques to be examined in a number of contexts; and the comparison of different approaches to constraint-based fuzzy qualitative simulation to be made.

This section contains an outline description of Mycroft which is based around the common division of the process (depicted in figure 1) into four parts: the model representation, the input data, the output behaviours and the inference engine. The model is a, possibly incomplete, representation of a physical system consisting of a number of variables and the relations between them. The variables of the system take values from a quantity space, which in this case consists of fuzzy values. The input data is the assignment of values to the exogenous and state variables, and the behaviours output are trees or graphs of states; each state being a consistent assignment of values to all the variables in the model. In each aspect of the Mycroft framework the representation of the model structure and the variables remains the same, while the inference engine utilised for the reasoning process changes. In the rest of this section each of these components is described in turn.

Fuzzy Sets and Quantity Spaces

Both qualitative reasoning and approximate reasoning have a common foundation, that of reasoning about system that are incompletely specified. However, the distinction remains, that whereas in qualitative reasoning a knowledge of the structure of the system under consideration is assumed, in fuzzy systems it is merely an



Figure 1: The Qualitative Reasoning Process

input/output representation that is utilised; albeit with the possibility of being empirically derived. Thus there are three advantages which ensue from the combination of fuzzy and qualitative approaches:

- the fact that the meaning of a qualitative value and its support set (the real number line here) are captured in a single representation,
- the ability to incorporate empirical knowledge into a model (which is also finer grained than the M^{+/-} constraint in QSIM (Kuipers 1986)), and
- being able to include more detailed knowledge of the temporal behaviour of the variables in a model than the total ordering available within QSIM, which is essential for use in such applications as model-based diagnosis and control.

This was the motivation behind the development of FuSim (Shen and Leitch 1993), which is the system which was the major influence on the development of Mycroft.

Fuzzy sets extend the ideas of traditional set theory to include the concept of partial (or graded) membership. It is assumed here that the ideas underlying fuzzy sets are known to the reader; however, a description of the domain and explanation of the concepts may be found in (Kosko 1992). In FuSim, for reasons of computational efficiency, trapezoidal fuzzy numbers and intervals are used. These are represented by the four tuple (a, b, α , β) which yields the interval shown in figure 2. The trapezoidal fuzzy numbers maintain the essential features of fuzzy sets: graded membership and the use of α -cuts.



Figure 2: A Trapezoidal Fuzzy Number

Scott 115

The quantity space which is built from fuzzy numbers must be closed, continuous, finite and cover all values which a variable can take. An example of such a quantity space is shown in figure 3. In fuzzy qualitative simulation, unlike QSIM, the quantity space for the derivatives of a variables may also be dense (that is can have any number of divisions consistent with the definition of a quantity space).



Figure 3: A Fuzzy Quantity Space

Model Representation

The Mycroft framework is a qualitative reasoning system within the so called constraint based ontology. The models used in the framework consist of sets of variables and the constraints that relate them. In fuzzy qualitative reasoning the operators utilised are the same as for its symbolic counterpart, though by the nature of the case there is a difference in the way they are implemented. The fuzzy case makes use of two concepts: the extension and approximation principles. The former of these extends the results of numerical and interval arithmetic and states that any operation on one or more fuzzy numbers will produce a fuzzy number as the result. All the variables of the system take their values from a predefined fuzzy quantity space. In performing any operation on one or more fuzzy values it is unlikely that the resulting fuzzy number will also be a member of the constrained variable's quantity space. Therefore, the latter principle provides a means of mapping the calculated value onto the appropriate quantity space. There are two ways in which this can be done: in the simplest case, if there is any overlap between a member of the quantity space and the calculated fuzzy number, then that quantity is a valid approximation and may be assigned as a valid value for the variable. This is depicted in 4. However, if there is more than one possible approximation (as is usually the case) then one can use a distance metric to gain a measure of the degree of approximation of each qualitative value. In this way one can prioritise the values from "most likely" to "least likely" (Leitch and Shen 1993).

In *Mycroft* the model constraints are causally ordered (Iwasaki 1988) and distributed over a number of differential planes (Wiegand 1991). That is, the qualitative differential equation (qde) model is developed on plane-0, and the relationships between the higher derivatives of the system are obtained by differentiating the qde and representing the results as a qde on the so called higher differential planes. To illustrate, consider the









single tank system of figure 5. The quantitative equations describing this system are:

plane-0

$$q_o = kV$$

$$V' = q_i - q_o$$

plane-1

$$\begin{array}{l} q_o' = kV' \\ V'' = q_i' - q_o' \end{array}$$

where q_i is the flow of fluid into the tank, q_o is the corresponding flow of fluid leaving the tank, V is volume of fluid in the tank at any particular time, and k is a parameter representing the resistance to flow presented by the outlet pipe.

This system model is linear and it can be seen that the relations in plane-1 have the same form as those in plane-0; with the difference being that each variable is the next derivative of the variable in plane-0. The means by which models represented in *Mycroft* can be used to generate the system behaviours is described in the following section.

Inference Engines

The subject of this paper is a novel synchronous fuzzy qualitative simulator developed within the *Mycroft* framework. As stated, the *Mycroft* framework contains a number of fuzzy qualitative reasoning algorithms; the majority of which operate asynchronously. In order to put the present work in its overall context this subsection will be devoted to a brief description of the asynchronous algorithms contained in *Mycroft*.

But first the distinction between synchronous and asynchronous operation must be clarified. Synchronous simulation is actually the more straightforward, since it is the most common approach to simulation in general, as manifested in numerical simulators. In synchronous dynamic simulation, the integration phase (Transition Analysis (TA) in qualitative reasoning) is driven by an external clock. That is, the length of the integration step is predetermined and proceeds relentlessly between the time-points specified. The solution of the model equations (Qualitative Analysis (QA) in qualitative reasoning) follows the integration step, and the simulation proceeds alternating between these two phases.

In contrast, pure qualitative simulators generally proceed asynchronously. That is to say, the length of time between the states is determined by the structure of the quantity spaces of the variables involved in the transitions. The transition analysis phase the values of the magnitudes and derivatives of, at least, the state variables of the system are known; and this information is used (along with transition rules) to decide which value these variables may take in the succeeding time interval (or time point). Thus no external driving force is required for asynchronous simulation.

Amongst asynchronous qualitative simulators there are two types of algorithm: constructive and nonconstructive. The classic example of the nonconstructive approach is QSIM (the basic algorithm of which is also utilised in FuSim). In non-constructive simulation the model of the system is used only to filter out assignments of values to variables. These possible values are assigned by the transition analysis phase on the basis of the knowledge of the present magnitudes and derivatives of each of the system variables individually. For someone acquinted with numerical simulation methods, this is counter-intuitive and appears to be a cause of spurious behaviour generation because values are generated which the constraints cannot eliminate. Therefore, Wiegand (1991) developed a qualitative simulator which sought to make constraint based qualitative simulation more constructive. In this approach the constraints of the system are used to calculate the values for all the variables of the system except for the exogenous variables and the magnitudes of the state variables. Thus, the problematic assignments of values to variables will not be made. However, regardless of whether the simulation is synchronous, asynchronous, constructive or non-constructive, the process is basically the same: it proceeds by alternating between the TA and QA phases.

The *Mycroft* framework contains three asynchronous algorithms: semi-constructive (which has a constructive QA phase and a non-constructive TA phase), constructive (which has a constructive QA and TA phase) and a fuzzy envisionment algorithm. Details of these are contained in (Coghill 1996). Each of these is designed to explore different aspects of constraint-based fuzzy qualitative reasoning. However, while temporal information can be calculated from each of these algorithms, when it comes to comparing the behaviours generated with real data, the temporal matching process is complicated. It is desirable to have a simulation engine which maintains the qualitative nature of the simulation while at the same time allowing the comparison of predicted and measured values at predefined time-points. Therefore the synchronous approach to fuzzy qualitative simulation described in the following section was developed.

Qualitative Euler Integration

Synchronous fuzzy qualitative simulation involves replacing the transition rules in Mycroft with an integration phase utilising a qualitative version of Euler's first order approximation to the Taylor Series: Qualitative Euler Integration (QEI). In the work presented in this paper the representation of our models is restricted to plane-0. This means that each state variable will have an associated qualitative magnitude and derivative whilst all non-state variables will only have an associated qualitative magnitude.

Given the time step Δt and a qualitative state, the Euler formula can be used to calculate an interval representing the predicted magnitude of state variables at the succeeding time point.

$f(n+1) = f(n) + f'(n)\Delta t$

This is made possible by our access to the alpha-cut of the qualitative values f(n) and f'(n). The alpha-cut of a fuzzy qualitative value is the crisp interval which represents those real numbers which have a membership value that is greater than or equal to a given alpha value $(\mu(x) \ge \alpha)$. Figure 6 shows a fuzzy qualitative quantity space Qm. (For the simulation α is set to 0.55.)



Figure 6: Fuzzy Qualitative Quantity Space Qm

To illustrate the use of QEI consider the following situation. Suppose the magnitude of a state variable X is defined over the quantity space Qm in figure 6 and the derivative of X is defined over the quantity space Qd (shown in figure 7. X is in the qualitative state [P-SMALL, P-MAX]. With an alpha value of 0.55 the alpha cut of P-SMALL in Qm is [0.155, 0.455] whilst the alpha-cut of P-MAX in Qd is [0.9551]. The predicted interval is given by euler integration, where Δt is chosen to be 0.35.

$$[0.503, 0.805] = [0.155, 0.455] + [0.955, 1] * 0.35$$

It is clear from this that at the new time point the magnitude of X will, by the approximation principle, be a member of the set{P - MEDIUM, P - MAX}. QEI is applied to each of the state variables and the elements of the power set of the results are passed in turn to the causal propagation phase.



Figure 7: Fuzzy Qualitative Quantity Space Qd

Continuity and Integration

In the previous section, QEI was introduced. However, even the simple example used there can illustrate the possibility that utilising QEI can produce discontinuous changes in the magnitude of a state variable. The state variable X could transition from P - SMALL to P - LARGE given the quantity spaces Qd and Qm and $\Delta t = 0.35$. Continuity, however, has been axiomatic in the development of all qualitative simulators including FuSim and Mycroft, and has effectively constrained the branching of partial envisionments. In this section, the maintenance of continuity in conjunction with QEI is described.

Under certain circumstances, the Euler integration step will only produce continuous changes in the magnitudes of state variables between time points. The circumstances under which this remains true depend on selecting the time step, Δt , such that the qualitative continuity condition for a single state variable is maintained. This condition is described below.

Consider the following situation; if Qm(i) is the quantity space for the magnitude of state variable *i* and Qd(i) is the quantity space for its derivative. We define the adjacent interval of a qualitative value v_i ($[v]^{adj}$) to be the interval representing the union of v with the adjacent value in the quantity space (which will be either greater or less than v depending on the sign of the derivative, d_i). That is,

$$[v]^{adj} = \begin{cases} v_i \cup v_{i-1} & if \quad d_i <_{\alpha} 0\\ v_i & if \quad d_i = 0\\ v_i \cup v_{i+1} & if \quad d_i >_{\alpha} 0 \end{cases}$$

For example, the adjacent interval of the qualitative value P - SMALL in the quantity space Qm is defined by the lower bound of the alpha cut of P - SMALL and the upper of the alpha cut of P - MED if the sign of the derivative is positive; and is defined by the lower bound of the alpha cut of ZERO and the upper bound of the alpha cut of P - SMALL if the sign of the derivative is negative.

Thus the we can express the continuity Condition for the QEI as follows:

Qualitative Continuity Condition

Select
$$\Delta t$$
 such that:
 $\forall a \in Qm(i) \quad \forall b \in Qd(i)$
 $a + b\Delta t \subset [a]^{adj}$

This condition must hold for every state variable i, in order to ensure that only continuous changes in qualitative variables are predicted between time points. Given the quantity spaces of our state variables we choose a time step to satisfy the qualitative continuity condition, in all situations.

Suppose, for example, that the quantity spaces of the magnitude of all state variables is Qm from figure 6 and the quantity spaces of the derivative of all state variables is Qd from figure 7. To find an upper bound on Δt we take the highest possible value in Qd, which is P-MAX, and a representative of the qualitative value with the narrowest adjacent interval in Qm, which in the present case is P - SMALL. The alpha cut of P - SMALL in Qm is [0.155, 0.455] whilst the adjacent interval is [-0.155, 0.745]. The alpha cut (where $\alpha = 0.55$ here) of P - MAX in Qd is [0.755, 1]. Substituting these values respectively into the Qualitative Continuity Condition gives the following inequality.

$[0.155, 0.455] + [0.755, 1] * \Delta t \in [0.155, 0.745]$

Here the derivative is positive, and so the largest possible value for Δt is that which yields 0.745 as the results of the QEI. Hence to ensure that the continuity condition is met:

$$\Delta t < 0.29 = \Delta t_{max}$$

Any Δt less than or equal to 0.29 will ensure a satisfaction of the qualitative continuity condition for the magnitude of all state variables in this example.

It should be pointed out here that this Continuity Condition has implications for the choice of quantity space that one may utilise. If one is merely interested in performing a simulation, then there is no problem; however, if external data has to be matched, then the quantity spaces must be constructed such that the upper bound for Δt is greater than the time interval between the data points.

Thus far the conditions for maintaining the continuity of the state variables throughout the simulation have been described. However, the question may still persist as to why one should want to apply such a restriction in the context of discrete time simulations. The justification is straightforward and is similar to that for the asynchronous case. However, since the same question arises in a broader sense in the context of applying the transition filters in the causal propagation phase, the justifaction will be left until after that process has been described. This is the subject of the next sub-section.

Continuity and Causal Propagation

In order to utilise the continuity filter applied in the causal propagation phase of Mycroft we must assume that the derivatives of the state variables and the magnitudes of all the non-state variables can only vary continuously over the given time step Δt . The assertion of continuity of the non-state variables and derivatives of state variables under QEI is equivalent to the use of continuity in the transition analysis phase in FuSim and the asynchronous algorithms in Mycroft. In these systems, since they both operate asynchronously, the temporal information is not used to calculate the next value a variable can take. Rather, the transition rules ensure that the next value a variable takes is one of the neighbours of its present value (and in the correct direction), and the times are calculated retrospectively. Therefore, the continuity rules are applied to all variables and their derivatives as matter of course since this is part of the definition of asynchronous simulation.

In the synchronous case, although the continuity of the magnitudes of the state variables is ensured by the continuity condition described in the previous section, this does not guarantee that the calculated successor values of the non-state variables will be continuous with their present values. Therefore the transition filters developed in *Mycroft* are used to ensure the continuity of the complete simulation and constrain the number of states generated.

In numerical discrete time simulation, continuity is taken as being implicit in the limit as $\Delta t \rightarrow 0$. In asynchronous simulation it has to be explicitly stated because all the information required to enable the simulation to proceed is embodied in the model representation. In the present case the question arises as to why one should not treat synchronous fuzzy qualitative simulation in the same manner as numerical simulation? The main answer to this question arises from the practicalities of the situation; namely the fact that the number of states generated would become unmanageable very quickly. Also, the states generated would be unconnected, making it impossible to ascertain which connections should be made to select a behaviour. An additional, and more conclusive, consideration in this regard is the fact that it has been shown that asynchronous qualitative simulators are conservative (Kuipers 1986); that is, they guarantee to find all the possible behaviours from the given starting state. Therefore if a synchronous simulator generates more states than an asynchronous one, then the additional states must be spurious.

As a supplement to this it may be noted that as Δt approaches Δt_{max} , it is not only guaranteed to transit continuously but also to do so maximally (that is, generate a value close to the furthest away boundary of the adjacent qualitative value). Therefore, values close to Δt_{max} will produce a behaviour tree close to or identical with those of the asynchronous case. Conversely, as Δt becomes smaller it becomes less likely that all the state variables will transit, assuming that they have different

values or quantity spaces.

Simulation of a Coupled Markets Model

In order to illustrate the foregoing and to assess the importance of continuity in QEI, a second order model of coupled Housing and Mortgage markets is constructed. This Coupled Markets model is a reduced form of a model which appeared in Wyatt, Leitch and Steele (1995). The two state variables are the housing price level P and the stock of mortgages M. The system contains two input variables, a market rate of interest r and a Housing supply Hs. The system is presented below as two differential equations. The first equation models the adjustment process of housing Price to clear the Housing market, whilst the second equation models the adjustment process of Mortgage Stocks to clear the mortgage market.

$$dP/dt = Hd1(P) + Hd2(M) - Hs$$
(1)

$$dM/dt = Md1(P) + Md2(r) - M \tag{2}$$

The quantity space for the magnitude of all variables is Qm from figure 6. The quantity space for the first derivative of the state variables is Qd from figure 7. The functional relations in the above equations (Hd1, Hd2, Md1, Md2) are modelled as degenerate fuzzy relations. Hd1 is a decreasing function of House Price P, since higher prices reduce demand for houses. Hd2 is an increasing function of mortgage stock, since the higher the mortgage stock the higher the demand for houses. Md1 is an increasing function of house price, since higher house prices lead to a demand for larger mortgages. Md2 is a decreasing function of the interest rate since higher exogenous interest rates lead to higher mortgage repayments. The final causally ordered model includes the degenerate functional relations and algebraic constraints to complete the Mycroft model of equations (1) and (2).

The partial envisionment for the coupled market model from an identical initial state was found by three simulation methods. A partial envisionment is a directed graph with a single root node (the initial state) in which the nodes are all the possible qualitative states reachable from the initial state. The first simulation method applied to the coupled markets model was the semi-constructive Mycroft described in (Coghill 1996). The second simulation method used QEI in the qualitative transition phase with $\Delta t = 0.2$ which ensured continuity in the transition of the magnitudes of state variables. In the second method, however, continuity filters were not applied during the causal propagation phase. The third method used QEI and applied continuity during the causal propagation phase.

Table 1 shows the number of nodes found in the partial envisionment for each method from the initial state:

$$\begin{split} M &= [N\text{-}SMALL, P\text{-}LARGE], \, Hs = [P\text{-}SMALL], \\ P &= [P\text{-}SMALL, N\text{-}MEDIUM], \, r = [N\text{-}SMALL]. \end{split}$$

The results in table 1 show that applying continuity filters during the causal propagation phase is important in producing a reliable qualitative simulator based on QEI. With the continuity checks and QEI 75 nodes were created in the partial envisionment. Without continuity checks 121 had been found and the partial envisionment was stopped whilst ongoing.

Simulator	S-C	Euler	E-C
Nodes	75	121 +	75
S-C: Semi-Constructive			
E-C: Euler with Continuity			

Table 1: Nodes in Partial Envisionment

Applying continuity filters during the causal propagation phase, therefore, leads to more tractable results from the QEI simulator. The extent to which continuity cuts down the number of possible transitions is clearly illustrated. The simulator without continuity filters in the causal propagation phase produces twice as many transitions as the simulator with continuity.

Table 1 also shows that QEI with continuity found the same number of states in the partial envisionment as the semi-constructive Mycroft, which is what one would expect from the argument of the previous section. Also, the results of the integration as used thus far have been treated as propagated values in the same manner as for the asynchronous case, and have been mapped straight back into the quantity space of the state variable by means of the approximation principle. However, this can, under certain circumstances cause problems. Consider the following case. A particular state variable has the magnitude value [P-SMALL P-SMALL] (from the quantity space Q_m ; and the time step Δt is 0.25. Performing QEI on this state variable yields the interval result [0.174, 0.53]. Now this maps onto the intervals for P-SMALL and P-MEDIUM from the quantity space. However, by the approximation principle, for all subsequent calculations the whole interval associated with the values P-SMALL and P-MEDIUM are used. This leads to a problem. In the case of P-SMALL, (if the same values of Δt is used then next integration step would yield the same values as before (an identical situation). In this case there are other values predicted by the integration step, but is possible, in the worst case, to select a sufficiently small value for Δt such that the calculated value does not reach the lower bound of the adjacent qualitative value; in this case the simulation would enter an infinite loop. A solution to this problem is presented in the next section.

Minimum Interval Euler Integration

The proposed solution to deal with persistence of qualitative states in QEI involves changing a representational primitive of the Mycroft system. The qualitative values of the magnitudes of the state variables will no longer have a single interpretation as an alpha-cut defined over the real line. Instead they will be used as a During integration the intervals calculated by the Euler method are recorded (instead of the alpha-cuts of the relevant quantity space) as the values of the magnitudes of the state variables. When integration is to be applied the next time around only these minimum intervals will be used. The problem of small intervals being approximated to the fixed quantity space during qualitative integration is therefore avoided. It may take a few integration steps for the minimum interval to traverse the alpha-cut of a qualitative value, thus causing that state to persist.

label for any sub-interval which is contained within the



Figure 8: Minimum Interval Euler Integration

Figure 8 illustrates such a process. Starting with a small propagated interval [-0.445, -0.339] (shown by dashed lines) and the qualitative derivative P-SMALL, the dotted line in the figure shows the result of applying Qualitative Euler Intergration with $\Delta t = 0.2$. Here, the variable cannot transit to the qualitative value ZERO in this integration phase because the interval [-0.445, -0.339] has not been approximated to the full alpha cut of N - SMALL. Assuming that the derivative remains P - SMALL throughout, the sequence of intervals the variable will have to transit through before reaching ZERO is show below:

[-0.445, -0.396], [-0.428, -0.337], [-0.411, -0.279][-0.394, -0.22], [-0.377, -0.161], [-0.36, -0.102]

Moving to a minimum interval representation for the magnitude of state variables addresses the persistence of qualitative states. This is because the definition of a node in the partial envisionment for minimum interval euler integration must take into account the minimum interval a qualitative value represents.

At the moment the use of the full alpha-cuts of qualitative values is retained in the casual propagation stage. This is because of the utilisation of degenerate fuzzy relations to represent empirical knowledge. This is an obvious problem to tackle to permit the extension of the idea of minimum interval integration to propagate only minimum intervals through the causal propagation phase. Thus there will be a complete coupling of the propagation methods of symbolic qualitative simulation and interval simulation. This will have the advantage of maintaining a consistent representation for all qualitative values whilst helping to eliminate the additional vagueness caused by the approximation of minimum intervals on the fixed quantity space.

Related Work

The work presented in this paper has focused on the use of an externally selected time step to synchronously drive the integration (transition analysis) phase of the simulation forward. Over the past few years there have been a number of systems developed which seek to make the temporal aspects of qualitative simulation more transparent; some of these have been synchronous and some have been asynchronous. In this section those most closely related to the present work are briefly reviewed for comparison.

The, deservedly, most famous research is that of Berleant and Kuipers (Berleant and Kuipers 1997): Q2 and Q3; and that of Kay and Kuipers (Kay 1996): SQsim. These systems are developed on and from QSIM, and are thus non-constructive. The primary simulation mechanism is QSIM and the quantitative information is then used to refine the simulation values and also the "time step" of the simulation. Thus the time is calculated as a result of generating future values and is asynchronous.

The constructive approach to qualitative simulation utilising semi-quantitative information has generated a fair bit of interest, and has often made use of fuzzy values. Unfortunately the trend in this line has been to converge on straightforward interval simulation. For example CA-EN (Bousson and Trave-Massuyes 1994) is a constructive qualitative simulator which may utilise either symbolic or interval values. CA-EN performs a fuzzy interval simulation which may be related back to an underlying fuzzy quantity space when required, though the actual simulation is basically interval based. It has led to the development of NIS (Vescovesi, Farquhar and Iwasaki 1995) which is an interval simulator.

Also in this category is the work of Bonarini and Bontempi (1994). This system is a fuzzy interval simulator (though it is still called qualitative). A great deal of useful attention is paid in their work to the issues arising from the interaction of the variables in complex dynamic systems.

In contrast to these systems it must be re-iterated that *Mycroft*, and the Euler integration approach contained within it, seeks to utilise the strengths of both synchronous simulation and qualitative reasoning, in order to enable the use of the simulator in contexts where the variable values may be expressed in vague terms but with a definite time stamp.

Discussion and Future Work

The ability to prioritise behaviours based upon their possibility with respect to the model may prove to be one of the major advantages of fuzzy qualitative simulation. This theme has been developed in (Leitch and Shen 1993). Current prioritisation has focused on FuSim where a distance metric is applied to the algebraic constraints in the model to identify the highest priority state, however, the transitions themselves are not prioritised. A prioritiser based upon the QEI should provide a more meaningful prioritisation of states because the transition of the state variable bears a causal relation to the next state which holds.

Another area which has yet to be explored is the generalisation of the Qualitative Continuity Condition to deal with models defined on higher differential planes. A higher order approximation of the Taylor series would be used during the qualitative integration phase. The more general Qualitative Continuity Condition would be defined over a time step Δt and the quantity spaces of a state variable on each differential plane.

Finally, one of the initial motivations for QEI is that it will allow the simulation of systems experiencing time varying and discrete inputs. Future applied work will involve the tuning of fuzzy qualitative dynamic models to Economic Time Series data where the economic system being modelled is subject to continuous exogenous shocks; as well as continuing the work on applying the technique in qualitative control and model-based diagnosis.

Conclusion

Synchronous Fuzzy Qualitative Simulation offers a complementary approach to asynchronous simulation with regard to the temporal aspects of the simulation. Instead of states being associated with overlapping temporal intervals, the use of QEI allows states to be associated with well defined time points. This has important implications for practical applications of qualitative simulation; simplifying the task of matching predictions against data and opening the way for simulation of systems which are experiencing time varying exogenous influences.

The effectiveness of simulators employing Qualitative Euler Integration has been shown to depend crucially on the application of continuity filters. Continuity of the magnitude of state variables can be ensured by the choice of quantity spaces and time steps which are consistent with the Qualitative Continuity Condition. Continuity for the highest derivative of each variable, however, must be imposed through the application of continuity filters during the causal propagation phase of Mycroft. Experimental results on a coupled mortgage and housing market model have shown that the application of continuity filters in the causal propagation phase produce more tractable simulation results than when they were not applied.

Comparison of the results of the semi-constructive Mycroft (Coghill 1996) and QEI with Continuity has shown significant differences in the handling of the temporal persistence of qualitative states. In particular QEI allowed infinite persistence times. This difference can be explained by the approximation of intervals propagated by QEI onto a fixed quantity space.

The proposed solution to handle temporal persistence of qualitative states under QEI involves maintaining the information provided by intervals propagated during the integration phase. The magnitude of state variables will be associated with the minimum interval of overlap between the propagated interval and the alpha-cut of a qualitative value in the magnitude quantity space of the state variable. By maintaining these minimum intervals for the next integration phase, minimum interval euler integration avoids the necessity of approximating small intervals on a fixed quantity space.

The final system trades some of the features of previous qualitative simulators for predictions at distinct time points. Quantity spaces must now be constrained by the Qualitative Continuity Condition. The semantics of the magnitudes of state variables must be adjusted to encompass any interval contained within the alpha cut of a given qualitative value. The exploration of these trade-offs have been made possible by the modular development framework provided by Mycroft; and that aspect of *Mycroft* reported here has been utilised in other research projects where models of multiple precision are required (Keller and Leitch 1994, Steele and Leitch 1997).

References

Berleant D.; and Kuipers B. J. 1997. Qualitative and Quantitative Simulation: Bridging the Gap. Artificial Intelligence 75: 215-255.

Bonarini A.; and Bontempi G. 1994. A Qualitative Simulation Approach for Fuzzy Dynamical Models. *ACM Transactions on Modelling and Computer Simulation* 4(4): 285–313.

Bousson K.; and Trave-Massuyes L. 1994. Putting more Numbers in the Qualitative Simulator CA-EN, in In Proceedings of the 2^{nd} International Conference on Intelligent Systems Engineering, 62–69, Hamburg, 5–9 September.

Coghill G. M. 1996. *Mycroft*: A Framework for Constraint-based Fuzzy Qualitative Reasoning, Ph.D. Thesis, Heriot-Watt Univ.

Fouche P.; and Kuipers B. J. 1992. Reasoning about Energy in Qualitative Simulation. *IEEE Transactions* on Systems, Man and Cybernetics. 22(1): 47-63.

Iwasaki I. 1988. Causal Ordering in a Mixed Structure. In Proceedings of AAAI-88, Pp: 313–318.

Kay H. 1996. SQsim: A Simulator for Imprecise ODE Models, Technical Report, TR-AI96-247, University of Texas Artificial Intelligence Laboratory.

Keller U.; and Leitch R. R. 1994. Fuzzy Qualitative Control, Research Memorandum, Heriot-Watt University.

Kosko B. 1992 Neural Networks and Fuzzy Systems, Prentice-Hall International.

Kuipers B. J. 1986, Qualitative Simulation. Artificial Intelligence. 29: 289-338.

Kuipers B. J.; Chui C.; Dalle Molle D. T.; and Throop D. R. 1991. Higher Order Derivative Constraints in Qualitative Simulation. *Artificial Intelligence* 51: 343– 379. Lee W.; and Kuipers B. J. 1988. Non-intersection trajectories in qualitative phase space: a global constraint for Qualitative Simulation. In Proceedings of the National Conference on Artificial Intelligence AAAI-88, 286-291, St. Paul, Minn.

Leitch R. R.; and Shen Q. 1993. Prioritising Behaviours in Qualitative Simulation. In Proceedings of the 13th International Joint Conference on Artificial Intelligence, Pp 1513 – 1528, Chambourcy, France.

Shen Q.; and Leitch R. R. 1993. Fuzzy Qualitative Simulation. *IEEE Transactions on Systems, Man and Cybernetics.* 23: 1038-1061.

Steele A. D. and Leitch R. R. 1997. Qualitative Parameter Identification, In Proceedings of the 11th International Workshop on Qualitative Reasoning, Pp 181 – 192, 1997.

Vescovesi M.; Farquhar A.; and Iwasaki Y. 1995. Numerical Interval Simulation: Combined Qualitative and Quantitative Simulation to Bound Behaviours of Non-monotonic Systems. In Proceedings of the 9th International Workshop on Qualitative Reasoning, 181– 190.

Wiegand M. E. 1991. Constructive Qualitative Simulation of Continuous Dynamic Systems. Ph.D. Thesis, Heriot-Watt Univ.

Wyatt G. J.; Leitch R. R.; and Steele A. D. 1995. Qualitative and Quantitative Simulation of Interacting Markets. *Decision Support Systems*. 15: 105–113.

Acknowledgements

This research was carried out while the authors were associated with the Intelligent Systems Laboratory, Heriot-Watt University. Edinburgh. One author (JAS) was funded by the ESRC to carry out a 1 year pilot project investigating the application of qualitative simulation techniques to Economics.

122 QR-98