Overview of modelling and analysis techniques for arbiters and related circuits

Fei Xia[*] and Alex Yakovlev[*]

1. Abstract

This is an attempt to give an overview of the state of affairs in the literature of the modelling and analysis techniques for arbiters and related flip-flop based circuits. Efforts are especially concentrated on the lower level modelling and analysis of simple circuits using analogue dynamic systems techniques. The phenomenon known as metastability is given particular attention, especially in conjunction with asynchronous operation of the circuits in question.

2. Introduction

Arbiters are circuits whose job is to grant, to its more than one clients, mutually exclusive access to a shared resource. Its use is wide spread among digital systems and circuits and it can be said safely that few such systems do not employ arbiters of one kind or another. Research interest in arbiters have been present in the literature for a long time, with a very rich body of results.

Metastability is a state wherein a normally bistable signal stays at an intermediate level between logic 1 (high) and logic 0 (low) for an indefinite period of time and appears to be stabilised at this level. In digital circuits it is usually the result of a system being put in one of its unstable equilibria and not coming out of it. According to dynamic systems theory such a state can persist for an unlimited length of time without outside triggering signals (such as noise). Classically, in such engineering fields as control systems and signal processing, it was assumed to be safe to disregard the existence of unstable equilibriums in systems as they were viewed to be unsustainable in the practical sense. Since the publication of actual observations of such states in arbiter circuits and especially because of the important implications of such states in these circuits, however, much research has been carried out to try and clarify the situation, both theoretically and practically.

Asynchronous operations in digital systems have always interested researchers because of their significance in such areas as real time, safety critical systems with a distributed nature. In such systems it is often advantageous to employ fully asynchronous processes. The complete elimination of time interference between concurrent processes makes it possible to accurately predict the temporal progress of each process in the system because the timing of each one is completely independent, not affected by the operations of other processes or the environment. This accurate prediction may be crucial if the system design must satisfy certain safety critical requirements. On the other hand, concurrent processes in a system usually communicate or otherwise interact with one another in some way and often share resources. Arbiters in charge of such shared resources therefore often need to be able to operate in the absence of synchronisation.

[*] Supported by EPSRC grant GR/K70715 (project HADES).

Arbiters operating without clocking or other synchronisation among its inputs (usually request signals from clients) theoretically cannot avoid being put into their unstable equilibria, and in practice have been observed to settle in non-trivial metastable modes of operation. This is an undesirable state of affairs as it may hold clients in a waiting state, thereby defeating the whole purpose of having an asynchronous design in the first place. Or worse still, metastability in certain arbiter signals may be interpreted by other parts of the circuit in different and potentially conflicting ways, thereby producing hazards.

Dynamic systems theory in its modern form can be traced back to the time of Newton. Mathematical tools available for the modelling, analysis and design of dynamic systems have progressed steadily since then, from the various forms of differential equations and their discrete counterparts, to the more concentrated and specialised classical stability theories, to the contemporary mathematical languages such as real and functional analyses and set theory. Such fields as control systems have seen extensive use of these mathematical tools. In circuit theory and systems these tools have also been widely employed. Ever since vacuum tube and semiconductor based circuits were first developed much effort has been made to develop convincing mathematical models for them. Digital and computer hardware systems have been modelled both on the higher, discrete, levels of operations and the lower, analogue levels to analyse their behaviours in various required levels of detail.

In this paper an effort is made to study the history and current state of affairs of the modelling and analysis techniques developed to tackle arbiters and related circuits. Of special interest are those involving analogue mathematics on the finer levels of system operation detail, especially for circuits and systems operating in an asynchronous environment and susceptible to metastability.

An attempt is made to trace one particular “family tree” of papers which especially interests the authors. Other published results are organised more or less around this core of work and based on the particular techniques employed.

3. Towards the ideal arbiter: the classical dynamic systems approach

In the simplest case, an arbiter has to deal with two clients who may request access to one shared resource. Thus such an arbiter must have a request input and a grant output for each of its clients. This is shown schematically in Figure 1.

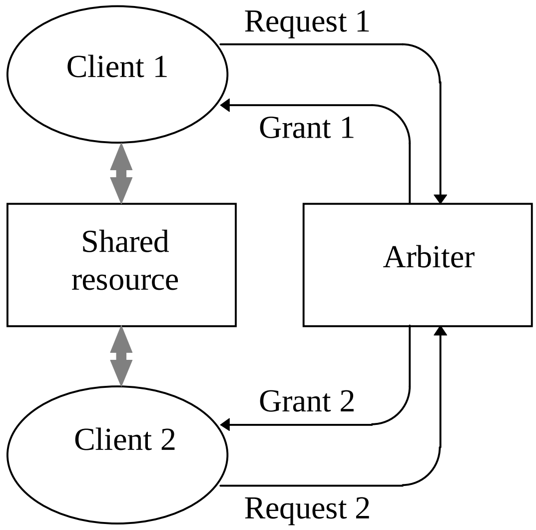

Figure 1 Arbiter with two clients

An ideal arbiter of this form must behave in certain ways so long as the clients conform to a set of behaviour rules. Intuitively, either client may raise its request, as long as its grant is low, at any time without regard to the states of the arbiter, the other client, and the shared resource; only when a grant is given may the corresponding client use the resource; and a client must lower its request (i.e. the request is withdrawn) once its use of the resource has been completed. Certain proofs also require that the clients be persistent, i.e., once a client has raised its request it must not withdraw it until it has received the corresponding grant.

When clients behave in the required fashion, an ideal arbiter is defined in [1] as follows:

- The arbiter is a passive element and is not allowed to initiate grants without the corresponding request having been raised.
- The arbiter must not allow both grants to be high at any one time. This is the mutual exclusion requirement without which circuits would not be arbiters. (called MUTEX in [1])
- After a request is withdrawn by a client, the arbiter must withdraw the corresponding grant within bounded time. This is to allow the next cycle of operations to begin. (RESET)
- Grants cannot be withdrawn until after the corresponding requests have been withdrawn. This gives clients control of how long they wish to use the resource. (DOMINANCE, constituting a HANDSHAKE protocol together with property 1 above)
- A grant must eventually be given if at least one request is raised. This is to say that a decision must eventually be made by the arbiter if either or both request signals are high. (LIVENESS)

Not really part of the arbiter, but more precisely part of the protocol, is the specific requirement that clients be allowed to raise their requests without any form of synchronisation with each other. Most importantly, concurrent requests must be catered for. (This final property is called CLOSURE)

It has been known for a long time that real life arbiters fail to perform according to the requirements listed above, or according to similar sets of requirements capturing the essence of all of the five properties. Specifically, the requirement that it must eventually make a decision and give a grant should requests persist, tend not to be satisfied under certain circumstances of client behaviour. This is a general characteristic of bistable circuits under conflicting inputs and not restricted to specifically designed arbiters alone [2-7] and was indeed at first most associated with the class of circuits related to arbiters known as synchronisers. The situation under which the circuits fail usually involve both requests (or the equivalent signals in synchronisers) being raised at or nearly at the same time, resulting in the circuits' inability to make a decision without significant delay. This irresolute state came to be known as the metastable state and much study was done to analyse system behaviour when it gets into metastability [8-11]. System failures result when an arbiter or related circuit hangs in the metastable state long enough (such as longer than an operation cycle of some other circuit connecting to it) to be propagated to other parts of the system and interpreted by these parts in different ways. Therefore on a practical level various efforts were made to minimise the onset or the effects of metastability [6-9, 12-14], but a claim of completely avoiding metastability by employing certain design tactics [15] was unfortunately shown to be faulty [16, 17].

Along this general line of research, several examples of work stand out as being most of interest from the point of view of modelling and analysis techniques employed.

The work that more or less generated true practical interests in the problem of synchroniser and arbiter failure owing to metastability was [3], where the authors were able to obtain photographic evidence of metastability lasting non-trivial periods of time on an oscilloscope. It became accepted that synchronisers and arbiters available then could not be guaranteed to produce a proper digital output signal within bounded time if they are subjected to asynchronous input signals.

The first significant work that as part of its overall scheme addressed the problem of metastability in digital circuits mathematically in a generalised manner was [18]. It was shown that a set of ordinary differential equations that has two stable equilibria to which most trajectories in the state space are asymptotically attracted to must have a region of indecisiveness. The conclusion is arrived at by way of the continuity argument: If the set of differential equation models is continuous, the state trajectories must also be continuous. If then there exist more than one stable equilibria their regions of attraction cannot overlap, and there must exist a non-empty set of states which do not belong to any such region of attraction between the borders of adjacent attraction regions. Figure 2 explains this argument with examples in one-dimensional state spaces.

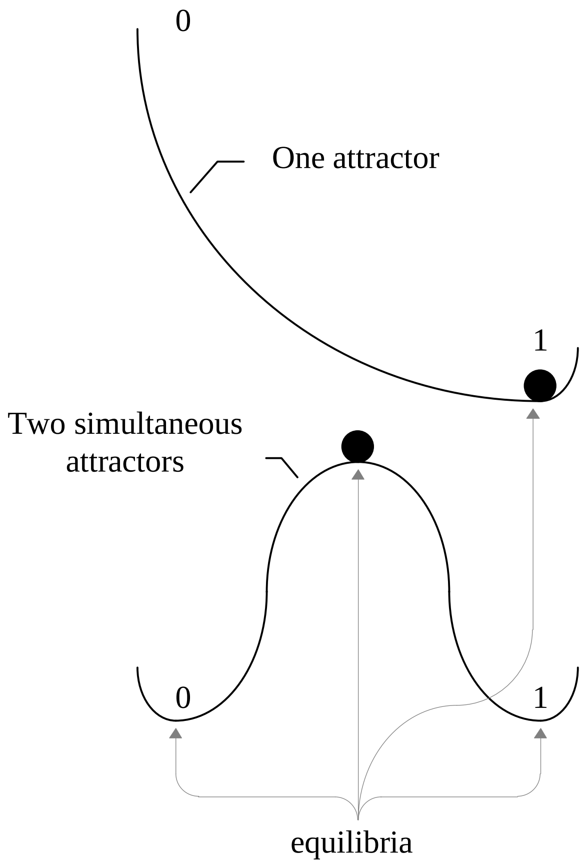

Figure 2 Stable and unstable equilibria.

The upper half of Figure 2 shows an illustrative mechanical system with a one dimensional state space when there is no conflicting input demands and the state signal has only one stable value — digital 1. The lower half shows the same system under conflicting input signals which created two stable equilibria, at digital 0 and 1. In this situation, assuming that the system is continuous both in time and signal level, the attraction regions of the two stable equilibria cannot overlap. This implies that there must exist at least one unstable equilibrium somewhere in between the attraction regions. If the state signal falls onto such an unstable equilibrium it cannot be mathematically proved that it will converge to one of the stable equilibria in bounded time.

The particular strength of the work of [18] lies in the fact that compared with most other work on arbiters and other bistable circuits it raises above any specific circuit design to attempt at obtaining generalised conclusions by analysing the behaviours of a class of mathematical models. If it could be argued that all bistable circuits can be accurately described by models like these the conclusions would then be extendible to the circuits. The main argument is that if digital bistable systems are assumed to operate on the Newtonian level, i.e., disregarding the arguments put forward in quantum mechanics and recognising signals as continuous both in time and in level, bistable digital circuits including arbiters and synchronisers are indeed quite likely systems that may be described by this type of mathematical models. The attempt of extending the conclusions to practical systems, however, ran into some difficulty as the analysis of [18] required certain strong assumptions (such as that inputs be fixed) on the models, and models of real bistable circuits could not be expected to satisfy them. It was not entirely clear whether this work had general relevance, or systems could indeed be designed that avoided metastability entirely.

Whilst most other research concentrate on the modelling, analysis and design of specific circuits using traditional techniques such as modelling with differential equations, computer simulations and laboratory observations, the work of [18] had great significance and was, indeed, the progenitor of most of the work studied below.

In [10], an attempt was made to improve on the results of [18] so that it's more applicable to a more practical field. This work and its improvement [11] do not have as strong assumptions as [18] and also made certain important extensions in the results. One such extension has to do with the topology of the region of indecisiveness. In [18] straightforward classical dynamic systems arguments were used such that the indecisive border region between areas of attraction is of a lower dimension than the state space. For instance, with two state variables the region of indecision constitutes curves, lines or isolated points in the state plane. It could be argued that the probability of a state falling into these border regions is effectively zero or infinitesimally small in a higher dimensional state space. No investigation was made to the possibility of states starting outside of the regions of indecision that nevertheless cannot arrive at a proper digital state within bounded time. The new results from [10] and [11] prove that even within the areas of attraction states near the region of indecisiveness may not reach the attractors enough to be recognised as having settled in a proper digital state within bounded time. This had the effect of enlarging the region of indecisiveness in bounded time to the same dimension as the entire state space. Therefore it is shown that the probability of metastable operation is nonzero. The models are argued to be general enough to cover all real digital systems.

As part of [19], an important theorem was proved such that the results of [10] and [11] may be extended from sequential machines to both combinational and sequential machines. This means that no Newtonian digital device can be expected to guarantee a proper digital output, given continuous input, within bounded time.

The techniques employed in [10], [11], [18] and [19] are very similar. Firstly, a mathematical model is established to describe the physical systems under discussion. Since such systems are not one specific circuit design but the whole class of bistable systems including arbiters, synchronisers and other flip-slop based circuits or even digital systems in general, the model must necessarily be that of a class of systems. Secondly, a way is found by arranging the models through algebraic exercises to allow the use of the basic continuity argument so that it is possible to show that under certain input signal assumptions metastability cannot be entirely avoided. This is a very powerful methodology although it can be said that the continuity argument itself had been well known and used in the field of dynamic systems for a long time. It was indeed part of any decent control theory undergraduate text book since the 1950s. Naturally the applicability of the models is always a problem and various convincing arguments are put forward to justify their specific construction. A very fine balancing act can sometimes be detected between making the models as general as possible so as to encompass all the possible realisations of the class of circuits under discussion, and arranging them in such ways so that the proofs can be derived. It is somewhat true that the more restrictions and assumptions are put to the system in the beginning the easier the proof but this results in more arguments as to whether the models can be assumed to represent real systems.

Another important property of the models constructed for these analyses is that they are in the classical dynamic systems sense. This is to say that a specific type of input-states-output relationship, with a state transition function in the form of

$$\varphi\colon \Sigma \times U \times R^{+} \rightarrow \Sigma$$

where Σ denotes the state space, U the input space, and R^{+} the time domain, is assumed for the system model. The mapping in the above formula consists normally a set of differential equations when the system is dynamic, as compared with a set of algebraic equations when the system is static. Specific properties associated with arbiters and similar digital circuits are provided by setting restrictions on function φ. Further arguments are centred on manipulating the conditions and the input signals to see whether metastability can be completely avoided by any system whose state transition function is of the required form which satisfies the restrictions.

This is the classical dynamic systems approach to such diverse problems as system stability and optimisation, and has been employed in for instance classical stability theory since the turn of the Century [20]. A shortcoming of this approach is that to someone not well versed in dynamic systems theory it may not be readily convincing that the models do represent the physical systems and that the restrictions on the models do not introduce unnoticed side effects. Also, by treating the system as a whole, the conclusion must be that metastability *somewhere* in the system cannot be avoided. This leaves the question open as to whether it is possible at all to construct circuits with the help of some kind of “metastability detector” or “filter” [21, 22] so that metastability does not occur at some crucial points in a system (such as the grant signals of arbiters). Indeed “metastable-free” arbiters have been reported whose implementation depend on these devices [14]. The author of [10] even stated later that flip-flops augmented with such detectors or filters can be free of ambiguous outputs [23].

The conclusion up to now can be summed up as follows:

- Digital systems that can be described by a set of differential equations that are continuous in state variables and time cannot completely avoid metastability in its state variables if given continuous inputs of certain types.
- In addition to the existence of initial states which do not produce a trajectory ending in a proper digital state, the probability of indecision within any bounded time is non-zero.
- If the task, such as asynchronous arbitration, is not required to be finished in bounded time, metastability-induced failure may not occur so long as certain design precautions are taken [14, 19, 21].

The ideal arbiter is still elusive.

4. A different approach to modelling and analysis

If the classical dynamic systems theory approach can be productive in the modelling and analysis of asynchronous digital systems including arbiters, why not try the more modern approaches that have come into vogue in such fields as control theory and stability theory [24]? Specifically, the method of setting out axioms in the beginning of a problem that describes the qualities of the systems being analysed in a qualitative, rather than quantitative, manner and by way of logical reasoning using such techniques as topology, set theory, real analysis and functional analysis [25-27] to arrive at desired conclusions would appear to be attractive. In fact, this approach would be especially attractive in such problems as certain discrete behaviours (such as the absence or presence of metastable operations) in digital systems which do not intuitively lend themselves readily to quantitative analysis. This is exactly what was done in the works listed below.

The equivalence and inter-realisability of three ideal hardware elements, the flip-flop (synchronisers and latches), the inertial delay, and the arbiter was demonstrated by way of setting out axioms that reflect the behaviours of these devices and then examining the logical consequences of these axioms [28, 29]. Although these works are not directly concerned with the metastability problem, the implication is that if one type of circuit cannot avoid metastability the other two cannot either. The axioms concerned describe input-output behaviours of the circuits in a qualitative manner and such classical dynamic systems theory techniques as differential equations are hardly made use of at all.

The work of [1] is based on the same technique. A set of axioms are set out for the ideal arbiter that encompass the qualitative descriptions of the properties CLOSURE, LIVENESS, DOMINANCE, RESET and MUTEX. It is claimed that the axioms are more relaxed than those found in [28] and [29], as time-boundedness is not now required for LIVENESS. This means that the arbiter would have to eventually give out one grant if at least one request is outstanding. Eventually implying that so long as the grant is given in finite time, all is well.

Again in order to facilitate the proof, the axioms are given in manners that on first look appear to be unnecessarily complicated. The proof resulted from a set of logical deductions from the axioms and it is shown that the four other axioms, taken together, are in conflict with MUTEX. In other words, any device that satisfies all of the other four axioms (two independent wires, each linking a request with its corresponding grant, is a trivial example of such a device), must fail MUTEX (as the trivial example surely does). The authors conjecture, without proof, that the problem is symmetrical, i.e. any device satisfying any of the four axioms must fail the fifth under certain circumstances.

Since the publication of [1] there have been some arguments as to whether the authors really did succeed in doing what they claimed to have done. Specifically, it has been argued that although they set out with the LIVENESS axiom assuming no upper bound for the settling time, later on in their proof this was erroneously changed to effectively introducing an upper bound for time. Such disputes have not been published, nor have them been answered by the authors in public.

Whatever significance the contributions the conclusions of [1] may or may not have, the novel employment of the “axiomatic technique” by these works cannot be overestimated. As it has solved, in more elegant and convincing manner, many qualitative problems in conventional fields such as control systems where dynamic systems theory have had strong applications, than more classical approaches, its potential in the modelling and analysis of arbiters and related circuits is demonstrated convincingly here as well.

More recent work in this direction include [30], where the techniques used in [1], [29], and [28], especially the axioms of [1], are employed in combination with the classical dynamic system approach of setting out differential equations, analysing the behaviour of the equations in the state space, and drawing conclusions from observations on state transitions. A couple of example arbiters are used to illustrate the possibility of proving that if a upper time bound is not required the initial states that do not have trajectories eventually converging on a proper digital state fall into a lower dimensional region than the whole state space. The probability argument is used to state that such arbiters “almost surely” will produce a correct response to any input. The author acknowledges, however, that carefully designed feedback control systems that use the arbiter’s states as feedback signals can cause the arbiter to operate in the indecisive region forever.

As a result of going back to the classical technique of setting out specific examples and then analysing them, the conclusions of this work is not as powerful as those of for instance [1] and [10]. As a further development of both to illustrate the difference between them, however, they are very illuminating. Now it can be said that if time bounded response is not required arbiters can be designed such that failure owing to metastability only occurs if the system falls into a set of states belonging to a region of a lower dimension than the whole state space. If it is accepted that the probability of this happening is infinitesimally small or zero as the author of [30] asserts, such devices can be used confidently in systems. If however bounded time is required for the response, system failure owing to metastability may occur for a region of the state space with an equal dimension to the whole state space. No convincing case can be made that the probability of this occurring being zero or virtually zero.

Probably of some significance, to date there has been no known laboratory observations of devices being kept in the metastable state for an arbitrarily long time.

5. Improving arbiters quantitatively

In parallel to the qualitative studies described above, much work has been done to find solutions to the problem of minimising the onset or effects of metastability in arbiter and related circuits. In general, the techniques employed in these efforts more or less fall into the classical dynamic systems approach. A system design is decided upon; mathematical models, usually in the form of differential equations, are established; analysis is made by way of state space studies, usually supported by computer simulations and sometimes by laboratory experiments; possible improvements to the design are suggested and further tested; comparisons with existing designs are made; etc.

Usually, it is considered sufficient to limit system models, specifically models of the semiconductor devices within the circuits, to a first order piece wise linear approximation of the physical system. Significantly, simulation with the help of such computer software as SPICE is often considered valuable and relevant in support of state space analysis and differential equation solutions [31-34].

Theoretical and practical techniques have been developed to reduce the incidence of failure in synchronisation hardware systematically to any given, nonzero, limit [35, 14, 21]. Much of the practical work employ classical dynamic systems theory techniques to some extent.

6. Achieving mutual exclusion without using arbiters

Is it possible to design a system where resource sharing does not require arbiters in the normal sense and still support fully asynchronous communications between processes? Various protocols and mechanisms have been proposed for communication between two asynchronous processes. The "slot" mechanisms described in [36-38] were devised for data communications between processes in the absence of any synchronisation. They employ multiple data storage slots, any one of which may be synchronised to either process at any time, but not simultaneously to both processes. The use of bit "control" variables makes the communication system globally asynchronous and locally synchronous [39]. In effect, the slot mechanisms realise "regular" and "atomic" registers in the data path between asynchronous concurrent processes with "safe" registers [40] to convey the values of bit-size control variables. This reduces the overall adverse effect of possible metastable operations in the data communication to a minimum, as the smallest data item that can be transmitted from one concurrent process to another is a bit. Significantly, it is claimed that even with the onset of metastability in its control variables causing conflicting interpretations of the signals in different parts of the system, the safety of the system is maintained and the performance would not be severely affected [38, 41], especially when metastability does not happen to the same signal repeatedly and would have settled when the signal affected is used.

However, bit variables transmitted asynchronously between concurrent processes usually require the use of some kind of latch (flip-flop) circuitry as conduits. Because of their memory flip-flops are related to arbiters in that they cannot avoid metastability all together. One implementation of the four-slot mechanism proposed in [37] is shown in Figure 3.

In Figure 3, the *d*'s are data slots (shared memory space where the data being passed between the reading and writing processes are temporarily stored) and *s*, *v*, *n*, *l*, and *r* are bit-size control variables. When writing or reading, a data slot is synchronised to the process in question and disconnected to the other process. The asynchronism occurs at the latch register circuits which maintain the control variables. This is how the slot mechanisms shift the asynchronism from the potentially large-sized data records to single bit variables, thus minimising the probability and effect of metastability.

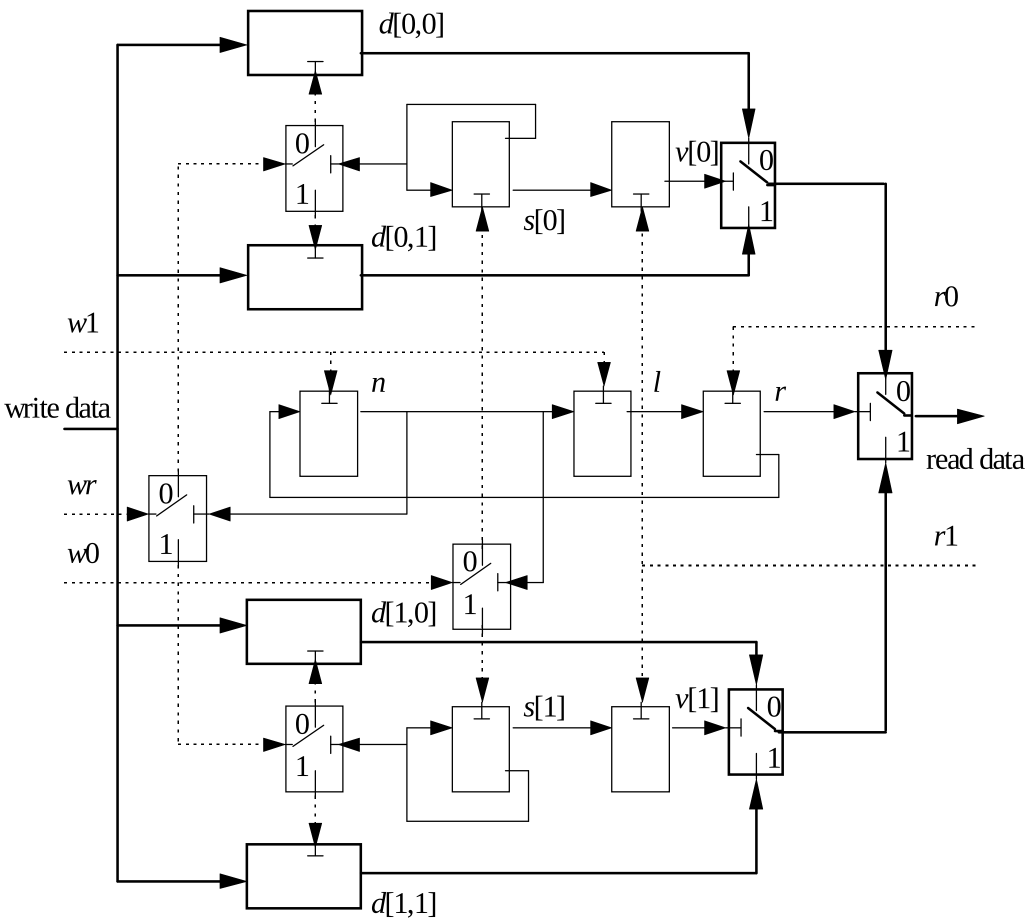

Figure 3 Four-slot mechanism.

From Figure 3 it is quite clear that devices prone to metastability, such as registers used in a completely asynchronous environment, can be connected to each other, making it possible for local metastability to affect other parts of the system in unexpected ways. The modelling and analysis techniques surveyed so far, concentrating on single devices and restricted to the analogue domain, are not very useful in system-wide analysis which might be needed for such systems as the slot mechanisms.

7. Concluding remarks and observations

Modelling and analysis techniques for arbiters and related circuit systems can be roughly grouped as follows:

- The quantitative study approach is based on classical dynamic systems theory and is normally supported by computer simulations and laboratory experiments. In this approach the objective usually is to find a better design such that the onset and/or effect of metastability is in some way minimised. Mathematical models of the systems are quantitative, normally in the form of differential equations. Solutions to the analysis problem usually involve either the solutions of the differential equations analytically or numerically or, if a trend needs to be established to guide design, analysis of the state space trajectories. Popular computer software used in such efforts include various versions of SPICE, and more lately, MATLAB.

- The qualitative study approach can be both based on more traditional dynamic systems methods or the more contemporary “axiomatic” methods. The goal of such analysis is usually to see whether a particular quality, such as the presence or absence of metastable operating modes, can be proved conclusively one way or another for a class of systems. In the former, more traditional approach systems are assumed to follow a set of differential equations and the solutions of such equations (the state space) are analysed in careful detail to arrive at the conclusions. In the latter a set of axioms describing the qualitative behaviours of the circuits in question are established and analysis proceeds with logical deductions from the axioms to try and establish a relationship among them. It is not cared whether the systems follow differential equations of a certain shape or not. In both of these approaches any computer or laboratory study must necessarily be a very illustrative nature and cannot be accepted as main supportive material.

Since most commercial computer hardware already contain arbitration and synchronisation circuitry of some sort and much of it do operate in asynchronous modes, it must be said that it is not without reason if someone regards the whole problem of metastability as not worth worrying about. However, on a more cautious and realistic level it is probably wise to regard metastability as a very real phenomenon whose probability of happening and especially causing serious failures in real systems, given careful design and testing, is very low. In addition, with the continued decrease of hardware clock cycles the issue of metastability will become more and more important.

The existence of system designs where metastable operations in some of its crucial parts do not cause failures in system operation must also be considered when designing asynchronous systems. Current modelling and analysis techniques in the analogue domain may not be very useful for systems with interlocking asynchronous devices which may go metastable, especially when the systems concerned are of some complexity.

There seem to be a number of areas that may merit future work:

- The better integration and cross support between quantitative and qualitative analytical techniques.
- The better understanding of the relationship between safety property failures and timing certainty and further development of the techniques that balance the trade-off between the two.
- The further investigation of more complex conflict resolving structures such as multi-way nacking.
- The development of an intermediate, discrete, modelling and analysis technique of metastability by representing the metastable state as a third level in addition to logic 0 and 1.

8. References

1 Mendler, M. and Stroup, T., Newtonian arbiters cannot be proven correct. *Formal Methods in System Design*, 1993(3):233-257, December 1993.

2 Catt, I., Time loss through gating of asynchronous logic signal impulses. *IEEE Transactions on Electronic Computers*, EC-15:108-111, February 1966.

3 Chaney, T.J. and Molnar, C.E., Anomalous behaviour of synchroniser and arbiter circuits, *IEEE Transactions on Computers*, C-22(3):421-422, April 1973.
4 Mayne, D., Minimise computer "crashes", *Electronic Design*, 9:168-172, April 1974.
5 Hurtado, M. and Elliott, D.L., Ambiguous behaviour of logic bistable systems, *Proceedings of the 13th Annual Allerton Conference on Circuit and Systems Theory*, pp.605-611, Urbana-Campaign, IL, USA, October 1-3, 1975.
6 Marino, L.R., The effect of asynchronous inputs on sequential network reliability, *IEEE Transactions on Computers*, C-26(11):1082-1090, November 1977.
7 Fleischhammer, W. and Dortak, O., The anomalous behaviour of flip-flops in synchroniser circuits, *IEEE Transactions on Computers*, C-28(3):273-276, March 1979.
8 Couranz, G.R. and Wann, D.F., Theoretical and experimental behaviour of synchronisers operating in the metastable region, *IEEE Transactions on Computers*, C-24(6):604-616, June 1975.
9 Liu, B. and Gallagher, N.C., On the metastable region of flip flop circuits, *Proceedings of the IEEE*, 65:581-583, April 1977.
10 Marino L.R., General theory of metastable operation, *IEEE Transactions on Computers*, C-30(2):107-115, February 1981.
11 Kleeman, L. and Cantoni, A., On the unavoidability of metastable behaviour in digital systems, *IEEE Transactions on Computers*, C-36(1):109-112, January 1987.
12 Elineau, G. and Wiesbeck, W., A new J-K flip-flop for synchronisers, *IEEE Transactions on Computers*, C-26(12):1277-1279, December 1977.
13 Calvo, J., Acha, J.I. and Valencia, M., Asynchronous modular arbiter, *IEEE Transactions on Computers*, C-35(1):67-70, January 1986.
14 Siddique, N. and Dike, C., Metastable-free arbitrator co-ordinates processors, *Electronic Design*, 14:107-112, April 1988.
15 Wormald, E.G., A note on synchroniser and interlock maloperation, *IEEE Transactions on Computers*, C-26(3):317-318, March 1977.
16 Chaney, T.J., Comments on "A note on synchroniser and interlock maloperation", *IEEE Transactions on Computers*, C-28(10):802-804, October 1979.
17 Wormald, E.G., Support for Chaney's "Comments on a note on synchroniser and interlock maloperation", *IEEE Transactions on Computers*, C-28(10):804, October 1979.
18 Hurtado, M., *Structure and Performance of Asymptotically Bistable Dynamical Systems*, PhD thesis, Sever Institute, Washington University, Saint Louis, MO, USA, 1975.
19 Chapiro, D.M., *Globally-Asynchronous Locally-Synchronous Systems*, PhD thesis, Stanford University, Report No. STAN-CS-84-1026, October 1984.
20 Liapunov, A.M. (1857-1918), *Stability of motion*, Academic Press, New York 1966.
21 Pechoucek, M., Anomalous response times of input synchronisers, *IEEE Transactions on Computers*, C-25(2):133-139, February 1976.
22 Kinniment, D.J. and Woods, J.V., Synchronisation and arbitration circuits in digital systems, *IEE Proceedings*, 123(10):961-966, October 1976.

23 Marino, L.R., *Principles of Computer Design*, Computer Science Press, Rockwell, 1986.
24 Bhatia, N.P. and Szergo, G.P. *Stability Theory of Dynamical Systems*, Springer-Verlag, Berlin, 1970.
25 Husain, T., *Topology and Maps,* Plenum Press, New York, 1977.
26 Bertle, R.G., *The Elements of Real Analysis*, Second Edition, John Wiley and Sons, New York, 1976.
27 Balakrishnan, A.V., *Applied Functional Analysis*, Springer-Verlag, New York, 1976.
28 Strom, B.J., Proof of the equivalent realisation of a time-bounded arbiter and a runt-free inertial delay, *Proceedings of the 6th Annual IEEE Symposium on Computer Architecture*, pp. 179-181, New York, 1979. IEEE.
29 Barros, J.C. and Johnson, B.W., Equivalence of the arbiter, the synchroniser, the latch and the inertial delay, *IEEE Transactions on Computers*, C-32(7):603-614, July 1983.
30 Mitchell, I.M., *Proving Newtonian Arbiters Correct, Almost Surely*, MSc thesis, The University of British Columbia, Canada, October 1996.
31 Kinniment, D.J., Yakovlev, A.V. and Gao, B., Metastable Behaviour in Arbiter Circuits, TR. 604, Department of Computing Science, University of Newcastle upon Tyne, UK, March 1997.
32 Kinnement, D.J., Yakovlev, A.V., Gao, B. and Xia, F., Towards Asynchronous A-D Conversion, accepted for publication in *Proceedings of Async'98,* San Diego, USA, 30 March - 2 April, 1998.
33 Kacprzak, T. and Albicki, A., Analysis of metastable operation in RS CMOS flip-flops, *IEEE Journal of Solid-State Circuits*, SC-22(1):57-64, February 1987.
34 Kim, L.-S. and Dutton, R.W., Metastability of CMOS latch/flip-flop, *IEEE Journal of Solid-State Circuits*, SC-25(4):942-951, August 1990.
35 Chaney, T.J., Measured flip-flop responses to marginal triggering, *IEEE Transactions on Computers*, C-32(12):1207-1209, December 1983.
36 Simpson, H.R., "Fully-asynchronous communication", IEE Colloquium on MASCOT in real-time systems, May 1987.
37 Simpson, H.R., "Four-slot fully asynchronous communication mechanism", IEE Proceedings, Vol. 137, Pt. E, No. 1, pp.17-30, January 1990.
38 Simpson, H.R., "New algorithms for asynchronous communication" and "Role model analysis of an asynchronous communication mechanism", IEE Proceedings on Computing and Digital Techniques, Vol. 144, No. 4, pp.227-240, July 1997.
39 Xia, F, Clark, I.G. and Davies, A.C., "Petri-net based investigation of synchronisation-free interprocess communication in shared-memory real-time systems", The Second UK Asynchronous Forum, University of Newcastle upon Tyne, Newcastle upon Tyne, UK, July 1-2, 1997.
40 Lamport L., "On interprocess communication", Parts I and II, Distributed Computing, Vol.1986, No.1, pp.77-101, Springer-Verlag, 1986.
41 Xia, F., Clark, I.G. and Davies A.C., Petri-net based investigation of synchronisation-free interprocess communication in shared-memory real-time systems, yet to be published, King's College London and University of Newcastle upon Tyne, UK, February 1997.