

## A Perspective on Concurrent Systems

(formerly MSc Lecture M2 by Meurig Beynon)

### 1. Varieties of Concurrent System

For the computer scientist, the term "concurrent system" typically evokes a context in which several processes are executing in a common environment. Each process resembles a sequential computation, in that it involves the reliable execution of a specified program of operations and interactions. The study of concurrent systems is concerned with finding ways to specify the patterns of interaction between processes, and to represent the intermediate system states. Specification languages such as CCS and CSP are designed for constructing precise and abstract mathematical models that represent the behaviour of concurrent systems of this nature.

Our primary concern in this course is with concurrent systems in a broader sense; systems that occur in everyday experience through the interaction of human agents, natural forces, and all manner of constructed artefacts, such as doors, clocks, cricket bats, timetables, computers, electronic and electro-mechanical devices. The study of such systems is motivated by many topical applications of computing: fly-by-wire, the intelligent car or house, virtual reality, customised environments for the disabled, computer-aided education, computer-assisted co-operative working.

Formal mathematical models of concurrency are not so directly applicable to complex concurrent systems of this nature. As Harel observes in [2], making comprehensible precise models of real-time embedded reactive systems is exceptionally challenging. The problems of representation and formal specification are especially difficult where the conception and design of such systems is concerned. In that context, an abstract stylised model of interaction is inappropriate, and the reliability and repeatability of processing and interaction patterns is hard to guarantee. It is significant that Harel refers to the need for visualisation (rather than simply a good mathematical formalism) to support reactive systems design [2]. Specification notations that have been designed for use in commercial concurrent systems design, such as generalised Petri Nets, statecharts, SDL, Lotos, have to make some compromise between expressive power and intelligibility, and precise operational semantics.

The motivation behind this course is twofold: practical techniques and theoretical principles.

- Can we develop more effective modelling techniques for requirements capture and design in modern applications of computing, where there is ever closer interaction between intelligent devices, people and their everyday environment?
- Can we identify fundamental principles for model construction, more general than those that a logicist tradition can deliver, to support more systematic development of complex concurrent systems?

The themes of this course reflect a more informal and human-centred perspective on concurrency than is represented in orthodox specification traditions. The archetype that informs our human perspective on concurrent systems is the concurrent system that we

experience through everyday life. Whereas a mathematical abstraction such as a CCS model has a public commonly agreed semantics and can be applied to make logically sound predictions about the behaviour of systems in a generic way, the concurrent system that we directly experience is subjective, located in the here and now, uncircumscribed and typically unpredictable. The principles of Empirical Modelling are arguably intimately connected with a commonsense perspective on concurrent systems that is derived from this experience. Key issues to be addressed in the course are:

- how do we conceive and represent commonsense concurrent system models?
- what conceptual stages are involved in understanding and taming concurrent systems to the point where we can invoke mathematical models?

In developing CSP [3], Hoare attached particular importance to the emphasis on processes and not on agency - hence Cooperating Sequential *Processes*. In order to construct formal models of concurrent systems with a precise operational semantics, it is much simpler to ignore agency, and to consider events rather than attributed events. In commonsense concurrent systems, where unpredictability is the norm, an unambiguous operational semantics is most inappropriate, and it is the *agent* rather than the process that is the primary concept.

## 2. Commonsense Concurrency

Concurrency as experienced is reflected in the dictionary definitions:

*concurrent*        running, coming, acting, or existing together  
*concurrence*     joint action, coincidence, assent

There is an implicit reference in these definitions to a present time and place. "Existing together" is meaningful relative to an observer, or at any rate mode of observation, that is external to the entities that are concurrent. Significant features of this observation include:

- subjectivity with respect to what exists together

What 'exists together' as of this moment in my mind is a very personal and loosely constrained conjunction of apprehensions of myself and my environment (the noise of a car, a pigeon on next door's roof, do I need my haircut?, an itch in the middle of my back, was that the door bell? etc. etc.) Even when we are in the same context, your concurrency and mine will be very different.

- discrimination in identifying observables and agents

'Existing together' presumes the identification of separate entities as well as their present coincidence. Concurrency entails both separations and conjunctions of experience. At a primitive level, my experience discloses only observables that are apprehended as having an integrity and identity (the pigeon's head, my hair, a tile on next door's roof, the colour of the sky) that may or may not be present currently. At a more sophisticated level, I apprehend characteristic clusters of observables whose presence is associated with particular sources of state change, and so can be construed as agents for this change.

- discretion over viewpoint and mode of observation

I have some degree of control over what I choose to regard as 'existing together' as of now. In concentrating on a task, I deliberately set out to ignore certain environmental influences, and to focus my attention on others. In so far as it's hard to direct attention at two things at the same time, I may also have a loose sense of how things are synchronised, and may regard things as existing together even though they can't be apprehended in one and the same instant. (I may also be unable to exercise discretion over some apprehensions: I can't ignore intense pain, and I can't distinguish the rapid but unsynchronised illumination of pixels in the screen display.)

In the context of this course, the role of the external observer as the mind within which the concurrency of a system is apprehended, or in some sense as the creator of a concurrent system, is fundamental. (Note that this is not to dissociate the external observer from the *effects* of agency - cf. the cricket watcher concussed by the ball.) The investigation of concurrent systems charted in this course is an extensive elaboration of the idea that the fundamental notions associated with concurrent systems (concurrency, agency, dependency, action, state) can be most appropriately formulated with reference to the interaction between a person (or person-like agent) and the world. Empirical Modelling is the means by which we seek to represent the external observer's conception of a concurrent system, as it evolves, typically incrementally, through experience of the system. This representation takes the form of an environment in which the observer can use visual metaphors, textual annotations, and analogue representations to imitate the interactions between components of a concurrent system as far as they are currently understood.

Formal models of concurrent systems are appropriate only in a context where the issues associated with subjectivity, discretion over viewpoint, and discrimination of essential entities have been resolved in an unambiguous way. The external observer in this case is an objective observer; the choice of observables and the characteristics of the agents are known to comply to the pattern required by the operational semantics of the model. A possible function of Empirical Modelling is to assist the empirical process by which appropriate selections of observables and agents can be identified. A more significant motivation for our radical shift of perspective is that formal models can only capture a

part of what is involved in the specification of more complex concurrent systems, such as reactive, social, or Virtual Reality systems. Empirical Modelling potentially offers powerful modelling principles and tools to guide specification of this nature.

### **3. Agency in Commonsense Concurrency**

Agency is a commonplace concept in our informal interactions with the world. Agency is associated with attributing state-changes to what is construed as their primary source. Its significance in everyday interaction highlights the limitations of formal models that factor out agency. Consider for instance: who's been eating my porridge? is there a doctor on board? this is the hole made by the bullet, keep away from the high radiation area, somebody might be listening ... all these are references to the relationship between state-changes and agents.

It is hard to reconcile the many contexts in which agency is invoked in commonsense concurrency with any single simple characterisation of the concept. The nature of the interaction between components of a system varies so much with context that it is difficult to say what components should be viewed as agents. Consider for instance the following four everyday concurrent system scenarios:

1. hearing a waterfall and birds singing
2. listening to a string quartet being played
3. observing the participants in a meeting
4. interacting with a word-processor

In 1, it may be reasonable to presume no communication, nor causal connection between the activity of the waterfall and the birds. If there is an interaction, it is hard to detect or represent it. It's not clear whether gravity should be invoked as an agent to describe the behaviour of the waterfall, or whether bird song should be attributed to a physiological source. Both waterfall and birds may still be treated as agents in a commonsense way: "that's the bird that makes a whooping sound", or "you can't hear its song properly because of the waterfall".

In 2, it might be convenient in modelling what is heard to regard the beating of time as the primary cause for the music. On the other hand, the playing is subject to variations from strict time, and inexact synchronisation, that can be attributed to the performers, or the composer. The music may also be subject to disruption because of an instrumental problem, such as the breaking of a string.

In 3, the participants have a clear autonomy, and their communication is explicit. It isn't clear to what extent the external observer has to take a view that accommodates their individual perspectives and private state changing responses. The status of the meeting as an agent is also an issue: some aspects of the state of the meeting (such as the current agenda item) are associated with the meeting rather than the individual participants; similarly, some state-changing actions are the prerogative of the meeting.

In 4, the status of the word-processor as an agent is obscure. At some level of abstraction, its state changes only through my actions; it doesn't do anything that I don't deliberately instruct, possibly even intend, it to do. On the other hand, when I forget how to perform a particular function, I may interact with it in an experimental manner. It can also serve its function erratically according to the context (e.g. it is unusable when sunlight falls on dust on the screen, one of the keys may cease to respond), or can serve a function beyond the scope of what it was designed to carry out (e.g. it can present a poem in an aesthetically pleasing way, it can give me a headache).

The difficulty of finding a good characterisation of the agent concept is highlighted by the range of issues raised by these four simple scenarios. To compound potential confusion, there are several meanings of agent in a dictionary definition:

*agent* a person or thing that acts or exerts power

any natural force acting on matter

one authorised or delegated to transact business for another

and, as cited by Luck and d'Inverno in [4], many different interpretations of the term "agent-oriented" in a computer science context.

The conventional way of improving our practical grasp on elusive concepts is to seek some form of mathematical or logical formalisation. In [4], Luck and d'Inverno aim to give a taxonomy for agents, together with a formal specification for each kind of agent. Formalisms have dubious relevance in the most creative phase of agent-based development (cf. Sloman [5]), yet it is in this phase that agency is most significant.

The Empirical Modelling framework to be developed in this course offers an alternative perspective on addressing agency in a principled manner. The key idea is that agency is in the mind of the external observer. It is shaped by the explanatory prejudices and requirements of the external observer, and by their past experience of the system. The model resulting from the Empirical Modelling process is essentially empirical and subjective, and reflects the pragmatic criteria that inform the explanation of system behaviour. A brief overview of the key concepts of Empirical Modelling is given in the next section. The detailed elaboration of the approach is the subject of the entire course.

#### **4. Analysing Commonsense Concurrency**

Our understanding of everyday concurrent systems is informed by past experience, present interaction and future expectation. Empirical Modelling addresses the process by which the external observer of a concurrent system develops such an understanding. It involves the systematic creation of an artefact, typically computer-based, that reflects the external observer's conception of the concurrent system. The artefact itself serves its function by metaphorically imitating the states of the system, and to some degree its significance is open to exploration in the same manner. Though this description suggests that the artefact has the character of a Virtual Reality representation, this need not be the case (nor can it be, with our current modelling tools!). The correspondence between concurrent system and artefact has more to do with conceptual similarity than

verisimilitude. By way of analogy, a librarian will find a library database much more useful than a VR representation.

The external observer can gain insight into a concurrent system in several ways. Observation of the system alone can provide a basis for insight. Cosmology, for instance, is entirely founded on observation of a system over which we apparently can exercise no direct influence. In other contexts, direct and directed intervention is possible: the external observer can interact with a system in the role of experimenter. Of equal importance in construing the behaviour of a system, the external observer may be able to carry out "off-line" experiments to test particular hypothesis about interactions within the system. In practice, the legacy of off-line experiments and previous experience of similar systems is very significant in determining the way in which a system is construed to operate. For instance, our knowledge of gravitational phenomena, gained from earthbound experimentation, shapes our cosmology.

The process by which the external observer develops insight into a concurrent system revolves around some key concepts: observables, current state, dependency and agency. Their significance and interrelationship will be sketchily described here, and elaborated later as appropriate. The discussion is deliberately framed in informal commonsense terms, with particular emphasis on the role played by empirical aspects of the external observer's activity.

#### **4.1. Observables and State**

Observables are the features of my environment that can be ascribed an identity. Their identification is a matter for empirical investigation: they typically come to be recognised over an extended period of interaction. "The pigeon on the TV aerial on the house next door": surely it's the same house as yesterday; think it's the same TV aerial; looks very much like the same pigeon. Observables can be arbitrarily abstract, so long as they have associated experimental procedures to determine whether they are currently accessible ("The pigeon has flown away"), and if so, determine their current status. Possible observables include such things as events: ("That's the 9.15 London train arriving"), quantities that are only indirectly measurable ("What's the current flowing through this circuit?"), and conditions whose direct apprehension is a matter of skill and experience ("This is tied with a reef knot").

The precise character of the observables in a system is determined by the motivation for study ("what I'm concentrating on") and the level of abstraction at which a system is analysed. The biologist and the ornithologist adopt quite different viewpoints on a pigeon, characterised by attention to distinct features of its nature and being. Observables must have some degree of persistence to acquire their identity, but typically come and go both out of current presence and existence. There are clear associations of observables into clusters that are typically coincident in respect of presence and existence: they form commonsense objects or entities.

In this connection, *presence* of observables is associated with what is commonly referred to as "current state of mind". A state of mind is similar to a real-world context, such as a physical location, in which the external observer can dwell ("I'm concentrating on doing

this at the moment"), which they can leave ("let's have a cup of coffee"), and to which they can return ("back to work now"). Continuity and change are perceived within a particular state of mind ("I've corrected the spelling in the first paragraph"): state change is associated with a change in value to a particular observable in such a context.

A further classification of observables will become an important focus of attention in this course. In a concurrent system, it is significant whether observables can be directly ("instantly") apprehended by the external observer or by agents within the system. Observables are often what we conceive other agents as directly apprehending: forces on wheels of a car, voltage across a light bulb. Equally significant is what we conceive an agent as capable of changing or responsible for changing: status of the handbrake, illumination in the room. Matters such as the identity, status and integrity of observables and agents are all empirically determined, and there are no absolute criteria by which judgements about such issues can be made. There are many ways in which this can be illustrated, and it is a useful exercise to the reader to devise commonsense scenarios in which identity, integrity, presence, existence, instant access etc etc are obscure or controversial (cf. Kent's *Data and Reality* for a bookfull of examples of this kind).

#### **4.2. Agency**

Commonsense agency is concerned with attributing state change. It is often appropriate to attribute state change to a commonsense object or entity ("the centre-forward scored the goal", "the wind blew the window open"). Semantic problems are encountered when we ask such questions as, which commonsense objects are agents?

A change in value to an observable can arise in a whole variety of ways. There is a sense in which a goal is scored by the referee, not the player (cf. Maradona - or was that the Hand of God?). Certainly, only the referee could possibly manipulate the score as arbitrarily as a fantasising football fan might. Perhaps a ball, goal posts and a goal line is also required? Most significantly of all, change in values to observables in a commonsense concurrent system can be associated with the external observer ("I didn't use to like that sentiment, but I'm comfortable with it now", "I should be looking at that in another way"). On this basis, it is difficult to discriminate between one commonsense object and another as agents. Indeed, even a primitive observable can have agency in this broad sense ("I was so pleased when I saw a tick against my answer").

Empirical Modelling reflects commonsense in viewing agents as having a different character according to previous experience and knowledge of a system, and the current context for its study. These views can be briefly represented as:

View 1: every observable or object is an agent, as is the external observer

Every observable / family of co-existent observables can potentially serve as an agent, being the cause, cue or trigger for some action of the part of another agent (be it only the modeller).

View 2: agents are objects responsible for particular state changes

Potential state changes may be invariably correlated with the presence of a particular agent (flag moves only if there's a wind, cell door can be opened if and only if it's been unlocked by the guard, train can depart only if the stationmaster

has authorised its departure). A stronger association between agent and state change is also possible ("I saw you do that").

View 3: virtual agency in the closed-world

The context for observation of a system is so circumscribed, the presumption of agent existence and presence is so secure, and the stimulus-response operation of agents is so reliable that an account of the behaviour of agents adds nothing to an account of the behaviour of the system ("Yours is not to reason why").

In the Empirical Modelling process, each of these views has a different status, and there is a tendency towards progression from the first to last views. View 3 reflects a stance that is often adopted by those who argue that agent-orientation is an unnecessary concept that self-evidently offers no advantage over a system view. Empirical Modelling is of interest somewhere between View 1, where the agent concept is vacuously broad, and View 3, where it's impotent. This interest centres around our uncertainty about the status of entities in respect of agency. In View 1, our concern is whether an entity exercises an influence over state changes a system; in View 3, our concern is whether all eventualities have been considered or encountered. Exploration of a system may lead us to reclassify an agent ("I've never seen it do that before", "there must be another agency at work", "I didn't know I could do that").

### **4.3. Interaction and Dependency**

The external observer is the archetypal agent in a concurrent system. Only the external observer can directly experience what it means to perform an action, or adopt a shift in perspective, that is truly a point of departure from what has been previously explored. Only the external observer can directly experience what it means to be surprised by an interaction.

A simplistic glance at history will perhaps help to put this concept of "surprise" in perspective. In Elizabethan times, it was apparently widely presumed that all that happened in the world could be attributed to one of several agents in a hierarchy that ranged from an omnipotent God to the lowest forms of life on earth (The Elizabethan Chain of Being). In such a framework, the task of the external observer is to attribute each state change to the agency commensurate with its importance (as in "Acts of God"). Expectations in this context were presumably far less well-developed than they are in our present scientifically and technologically oriented culture. Unpredictability and unreliability were characteristic of most experience, and the View 3 perspective on agency would doubtless have been far less well-populated than it appears to be in our time.

From this discussion, it is clear that our commonsense concept of agency, and our capacity to be surprised, is closely connected with the explanatory framework within which we operate. For the scientific reductionist, it is utterly implausible that a Rolls Royce should give birth to a Mini. Agency is the antidote to surprise ("There are more things in Heaven and Earth that are known of in your philosophy", "With God, all things are possible"). To the rationalist, agency is a recourse that is used where expectations are confounded: an acknowledgement of ignorance, and incomplete understanding. Russell [6] approved of relativistic explanations precisely because they dispensed with agency.



As participants in everyday concurrent systems, we all have experience of agency outside the scope of View 3. Agency changes its character totally only when we become so convinced of the outcome of interactions that we regard them as incapable of surprising us. The very same shift of perspective underlies the paradoxical way in which we use the term *experiment* for activity in which the outcome is uncertain, and for activity where we would be amazed if our expectations were to be unrealised. And, indeed, we use the term in both senses to refer to the very same activity.

Typical interaction with everyday concurrent systems combines activities that confirm expectation with exploratory activities. The apprehension of agency is associated with first recognising the extent of our own influence over our environment, then observing state changes outside the scope of our control. This classification of experience is subject to empirical validation (when does a baby learn to recognise that it can control its own image in the mirror?, when do we become aware that someone has been secretly accessing our computer account?). Agency first manifests itself in the transition from a View 1 to a View 2 status (it becomes apparent that no one accesses my account when my sister's boyfriend is on holiday).

Recognising the extent of our influence over our environment is closely related with the identification of dependencies between observables. Dennet [1] p428 describes "the basic method of obtaining self-knowledge ... about our own internal states, tendencies, decisions, strengths and weaknesses" in these terms: *Do something and "look" to see what "moves"*. That this plays a fundamental role in brain function is evident from experiments on chimpanzees that Dennet cites:

A chimpanzee can readily learn to reach through a hole in the wall of its cage for bananas, guiding its arm movements by watching its own arm on a closed circuit TV camera mounted quite some distance from his arm ... a decidedly non-trivial bit of self-recognition, depending as it does on noticing the consonance of the seen arm movements on the screen with the unseen *but intended* arm movements.

It is interaction of this general nature between the external observer and a concurrent system that Empirical Modelling is intended to represent. In this context, there are certain presumptions about the external observer that are quite characteristic of human (or at any rate human-like) agency. Some capability for interaction with the environment ("Do something"), a perception of state-change, and of identity ("look to see what moves"), combined with ability to remember and so to have an expectation, and to correlate the action with what moves. The precedents for human-computer interaction that exhibits these characteristics are to be found in the spreadsheet: the development of our perspective on multi-agent systems originates with a close study of 1-agent systems in which the paradigm for interaction is, at a conceptual level, similar to setting up and using a spreadsheet.

## **Summary**

The main moral of this introductory account of commonsense concurrency and Empirical Modelling is that interaction is central to the notion of concurrent system, and that concurrency, agency, dependency, action, state are notions that can most appropriately be formulated with reference to the interaction between a person (or person-like intelligent

agent) and the world. The discussion has touched on a number of important themes to be revisited in subsequent lectures (see Empirical Modelling for the Single Agent and Concurrent Systems Modelling: Agentification, Artefacts, Animation.)

### **Concluding Remarks**

The ideas developed in this discussion of commonsense concurrency are so familiar - even commonplace - that their relevance to any practical study of computer-based modelling seems questionable. The significance of Empirical Modelling in this context is that it is concerned with the representation of commonsense concurrency - an issue that - it will be argued - cannot be satisfactorily addressed by traditional formal methods. Traditional methods of specification are irredeemably tied to closed worlds, and even if these closed worlds should be extensible, they are only inhabited by the agents of View 3. In contrast, the environments of Empirical Modelling enable agents associated with all three perspectives to co-exist, just as they do in our everyday concurrent systems (even Solaris sometimes crashes, even now I can still discover things I didn't know I could do with Unix). Our objective in such modelling is to represent the reliable and predictable operation of a View 3 agent in such a way that it is still possible to override automatic operation, as in simulating malfunction, and to be able to accommodate the realisation that there are latent dependencies between View 1 agents that have hitherto been unremarked.

### **References**

1. Daniel C Dennett *Consciousness Explained* Penguin Books 1991
2. David Harel *Biting the Silver Bullet* IEEE Computer January 1992, p8-20
3. Tony Hoare *Communicating Sequential Processes* Prentice-Hall International
4. Michael Luck and Mark d'Inverno *Agency and Autonomy: a Formal Framework* CSRR#276, Univ. of Warwick, Dec. 1994
5. Aaron Sloman, Brian Logan, in Proc. FOMAS Workshop Warwick 1996
6. Bertrand Russell *The ABC of Relativity* George Allen and Unwin 1971