WORKING NOTE AC107

Alan Creak
9 December 1996

KITES AGAIN, WITH MORE VTA.

Not a very long time ago, I tried to use the VTA (vocabulary translation analysis) method to analyse the control of steerable kites using a single string[1]. From my point of view, the result was a minor triumph, for I was able to infer the principles of the system from the behaviour despite having no detailed knowledge of the nature of the kite nor of its dynamics. My subsequent guess at a mechanism was close in principle, but not quite right in detail[2]. In this note, I offer what I think is a better treatment, expressed in terms of the significantly more evolved version of VTA which has grown, without any proper documentation, over the past few months, and incorporating just a little of the correct mechanism.

I had guessed that the signal on the control line must be explicitly encoded in terms of *changes* in tension, and that the kite must be so constructed that (for example) sudden jerks switch it between two or more stable states, so that they affect its response to more gentle changes over an extended period. It turns out that the true situation is slightly different. Most of the details inferred from the VTA analysis are correct, but the state mechanism is not as I had imagined it. Instead of selecting a stable geometric state by a pulse code, the string tension switches the kite between three states[4]. With a slack line, the kite is planar and tumbles; with a tight line, the kite is distorted from the planar into a stable configuration in which it advances in the direction in which it is currently pointing; while at an intermediate tension there is an equilibrium position where the kite rides in a stationary position. The experienced kite pilot uses the visual feedback from the kite's behaviour to determine what sort of movement in what direction is desirable, then manipulates the line tension to rotate the kite to the desired direction and drive it to the required position.

As Geoff pointed out[2], the kite's state change is gradual. The string tension is used as the medium for a *fuzzy language*, with the kite changing continuously from one state to another as the tension is altered. This fuzziness doesn't seem to make much, if any, difference to the VTA model, so I've ignored it.

It is interesting that the system description derived by starting with a more complete knowledge of how the kite works differs hardly at all from that of my earlier attempt, where I was trying, without knowing the mechanism, to invent a model which would behave in a prescribed way. The descriptions would differ materially at a deeper level if analysed with greater attention to the mechanical behaviour and aerodynamics, but I think that the success of the VTA approach shows that working at the level of information transfer can give useful results even without knowing the details of the system.

SYSTEM DESCRIPTION, AND SOME NOTATION.

The system has three physical components, the pilot, the string, and the kite. The interactions between these components are illustrated in the flow diagram below.

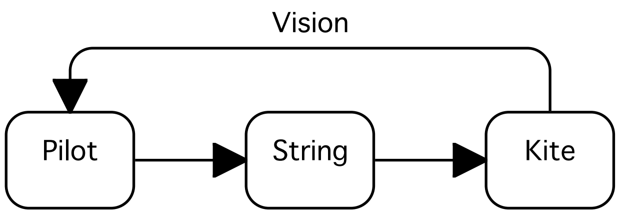

This is not a very exciting flow diagram, but it does show the essential nature of the interactions as a feedback loop.

The components are very different in their degrees of complexity. The pilot can operate at a high cognitive level, observing the kite's position, planning its future path, and translating these plans into manipulations of the string. The string can transmit these manipulations to the kite, and the kite can "interpret" the signals it receives and express them in terms of its internal configuration, but that's all.

The consequent behaviour of the kite is, oddly enough, not really something which the kite "does". All the kite "does" is maintain its state; the ensuing motion is not determined by any action of the kite. Instead, it is caused by interactions between the kite and environmental phenomena. The most important external influences are the wind, which exerts a force on the kite determined by its state, and time, through which the force determines the kite's position. I have, less than enthusiastically, regarded this part of the behaviour as a property of the physical frame in which the behaviour occurs, because I can't find anywhere else to put it.

It is convenient to show the relationships between levels of complexity and system components in a *cell diagram*, in which a cell is assigned to each system component at each complexity level significant for its operation. Here is a cell diagram showing the relationships :

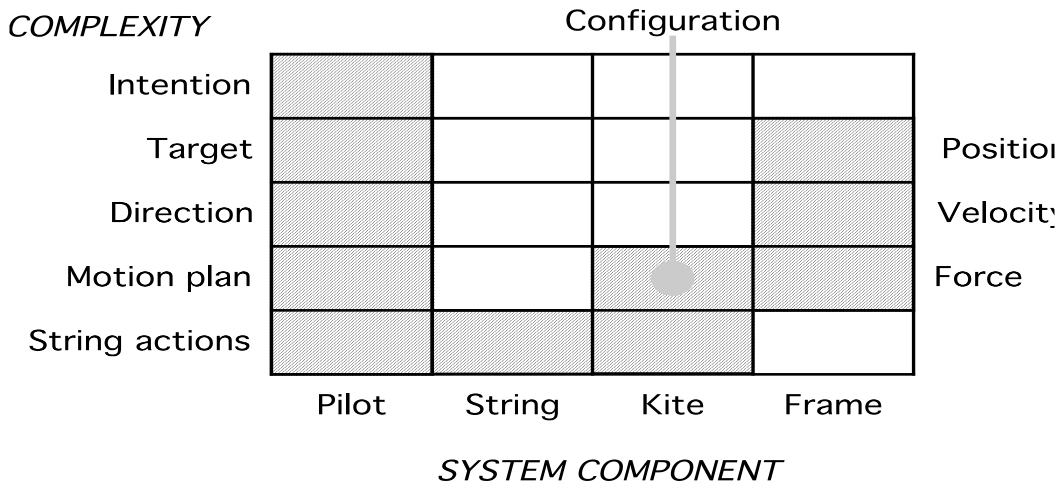

This diagram is, in effect, a rudimentary and informal device for taking into account Ashby's *law of requisite variety*[3]. According to this law, a mechanism dealing with a system of high complexity must itself be correspondingly complex; therefore, if the string and kite are to control information of comparatively high complexity, it must be encoded in terms of the information which the simpler components *can* carry. In other words, if the medium cannot transmit raw signals of the required complexity, a vocabulary which can be carried by the medium and which *can* encode such signals must be devised.

While the cell diagram shows what sorts of path through the system are available, it does not give details of vocabularies which must be used. The vocabularies used along a possible communications path, cell by cell, can be presented – or, at least, hinted at – in a *vocabulary table* :

#	*inf*	*Form*	*Locus*	*Using*	*Medium*	*Vocabulary*
1		Intention	Pilot		?	Thoughts
2	→	Target	Pilot		Coordinates	Spatial positions
3	→	Direction	Pilot	10	2-D vector	Spatial directions
4	→	Motion plan	Pilot		Kite motions	Kite actions
5	→	String actions	Pilot		Hand movement	{ Pull; hold; relax }
6	→	String signal	String		String tension; time	{ High; equilibrium; low }; any interval
7	→	Kite state	Kite	Time	Geometry; orientation	Set of kite states
8	→	Kite forces	Frame	Wind forces	Forces	Force vectors
9	→	Kite velocity	Frame		Velocities	Velocity vectors
10	→	Kite position	Frame	Time	Coordinates	Spatial positions

This table is primarily intended to show the steps of a single communication path, though with ingenuity it can be stretched a bit to take into account loops and branching.

- The first two columns (headed # and *inf*) number the steps and give some indication of the primary source of information for the step – usually the preceding step, indicated by a right arrow.

The blank in the first row leaves the source of information unspecified; a link to information from a different step can be denoted by the step number.

- The third column (*Form*) shows the nature of the information represented by the step as a brief description of its significance in the system.

- The fourth column (*Locus*) identifies a system component as the situation of the information. In most cases, the information has some more or less identifiable physical manifestation which can be identified within a component. As I pointed out above, in this example the "real" dynamic properties of the kite (as opposed to the pilot's perceptions of the properties) are not so easy to categorise, so I've left them as properties of the environment at large.

- The fifth column (*Using*) lists any external influences other than the primary source of information which are significant in determining what happens in the step; a number in this column refers to the information identified in the corresponding row of the table, while external phenomena are described in words.

- The sixth column (*Medium*) identifies the medium in which the value is expressed. This need not be physical; it is whatever changeable set of parameters are used to encode the quantity to be conveyed.

- The seventh, and final, column (*Vocabulary*) defines the vocabulary used as a subset of possible values, or ranges of values, of the disposable parameters of the medium.

REMARKS ON USING THE NOTATION.

The columns headed #, *inf*, and *Using* make their first appearance here. They are all to do with showing more clearly whence the information used by a step comes. The result is to extend the table beyond a list of possibilities to one which contains more, and more orderly, information about the system.

As is perhaps clear from the diagrams presented, I do not think it necessary to insist that the method be used with any extreme rigour. This is quite important, because in many cases – within people's brains, when dealing with fuzzily defined languages, etc. – it would not be possible to identify loci or vocabularies for the different signals with any degree of certainty, and indeed the signals themselves might be hard to find by experiment.

It is far more important that the resulting description be plausible, because then it can be checked. If it should later turn out that the predicted behaviour does not coincide with that observed, then some revision of the suggested model will be necessary, but by that time it is likely that the exercise of drawing up the first model will have led to much deeper understanding of the system that at the start. This is certainly a good description of the progress from the first to the second (this) model of the kite system.

A specific example is seen in the vocabulary table, which shows the kite's velocity as derived directly from the wind forces on the kite. This is not even qualitative physics; it could perhaps be called intuitive physics. In a more detailed physical analysis, the velocity of the kite given a particular system configuration is determined by the interaction of gravitational and aerodynamic forces acting on the kite, and equations of motion must be solved to determine the trajectory. Using VTA, it is sufficient to know that in fact increasing the string tension causes the kite to move with some velocity, and the table shows that relationship. Of course, if the fact turns out to be wrong, then the analysis fails.

The use of the "frame" as a repository for awkward quantities is not entirely satisfactory, but it's the best I can do at the moment. What should be done with steps in the representation of the information which seem to have no real physical embodiment within any of the system components ? It is not unreasonable to suppose that the pilot's notions of his intention in flying the kite, and his immediate plans for the next control action, and the other steps which I've listed, should have some sort of detectable representation within the pilot; likewise, the string tension lives within the string, and the kite state is determined by the current geometry of the kite. But there's nothing in the kite which can show its velocity or position. It might not be accidental that terms like "velocity" and "position" are abstractions, with no obvious real physical counterparts. I suspect that I might have to withstand an attack of philosophy if I go much further with this argument, so I'll stop, but any comments would be of interest.

The appearance of several "layers of complexity" in the cell diagram for the "frame" begs questions. It is by no means clear that the position of the kite is any more "complex" than the force on the kite, even though there are slightly suggestive similarities between some of the items in the first and last columns. Perhaps the attempt to allocate corresponding levels of complexity across the diagram is misguided, though it seemed to work well enough for communication aids. A difference is that the entities in the *Pilot* locus are human percepts, and we can easily see the process from step 1 to step 5 as one of reducing an originally complex intention into simpler and simpler terms, until the result is simple enough to be given to a machine, while there is no one perceiving the entries in the *Frame* locus.

REMARKS ON THIS MODEL.

This model is very like that which I suggested before[1] – indeed, so far as the descriptions go, the two are practically identical. The extended notation used this time gives a much more thorough account of what's happening, and the cyclic nature of the interaction is much clearer.

That the two descriptions are nevertheless so similar shows that the detailed behaviour of the kite in its different states is not described, for that is where the two models of kite behaviour differ. In my first model, I assumed that a kite could be switched by coded signals on the string between three different, and essentially stable, configurations with different aerodynamic behaviour. In fact[4], the signals are even simpler than I had supposed, with a change in tension switching the kite from its normal equilibrium state into one of two possible mobile states, in one of which the kite changes its orientation while in the other it moves bodily. Should this difference be apparent in the descriptions ?

I think that the answer is "only to the extent that it's useful". In this description, the details are not particularly important, and are included in the determination of the kite motion in steps 8 to 10 of the vocabulary diagram. If more detail were required, these steps could be expanded to show two different behaviours corresponding to the two states. It makes the table more complicated (it seems to be easier in practice to deal with such branches by setting up separate tables for the two cases), but the flow diagram would clarify the picture.

REFERENCES.

1 : G.A. Creak : *Kites : introducing HKI*, unpublished Working Note AC99 (July, 1996).

2 : G.C. Crumplin : various conversations, June to November, 1996.

3 : W.R. Ashby : *An introduction to cybernetics* (Wiley, 1958).

4 : G.C. Crumplin : *Not an Indian fighter kite* (Private publication, 1995).