Alan Creak
9 July 1996

KITES : INTRODUCING HKI

Geoff Crumplin has devised a steerable kite with only one control line. I have a way of describing interactions which I believe to be in principle universal. While brooding on Geoff's claim, it occurred to me that I should be able to treat it as a problem of control in terms of my description. Here I try to apply my technique to the question of flying kites.

INTRODUCTORY REMARKS.

My knowledge of kites is as close to zero as makes no difference. I've seen people fly kites, watched sailing boats occasionally, and stood in the wind. My knowledge of human-computer interaction is somewhat, though not a lot, more extensive, and that, together with my interest in rehabilitation systems, has led me to devise a design aid[1] for such systems. In my approach to design, strong emphasis is placed on the nature of the message intended to be transmitted through the system and its encoding in terms of physical variables in different parts of the system.

Geoff Crumplin mentioned in a conversation[2] that he had designed a kite which could be fully controlled by a single line. To be precise, I should qualify that in case I'm misrepresenting him : I think that's the substance of what he said, but being less than obsessed by kites I might have misunderstood. In any case, I found that the thought of steering a kite with one string kept on nagging at my mind from time to time as an unresolved problem. It seemed to me that the information-carrying capacity of the single string must be to be too small for anything but the simplest control, and I couldn't imagine how it worked. Geoff's brief explanation that it was a matter of the kite's distorting asymmetrically (subject to the same proviso as before) did not seem to help much.

Eventually, it occurred to me to think of it in terms of my interface approach. It is well suited to the task : the set of instructions to be transmitted to the kite seems to be small (though at least some of them include real values, which is a complication which I hadn't addressed before), and the communication channel is rather simple. What follows is the result of my thinking.

A final remark is appropriate, if only to forestall potential objections. It will not surprise me at all if much of what I have written is wrong. I've already registered my incompetence in kiting, and will later similarly protest my incompetence in aerodynamics. If you identify such errors, I'll be interested to hear about them. Even so, it doesn't matter a lot, because my primary intention is not to fly interesting varieties of kite; it's to test my design aid on a different sort of system, and at that level I think it's worked rather well. It has given me a framework within which I have been able to analyse an unfamiliar problem, and it has focused my attention on the signals which must be communicated to the kite and the corresponding demands on the communication channel. And that's what I wanted it to do.

THE SYSTEM.

I first imagined the complete system as looking something like this :

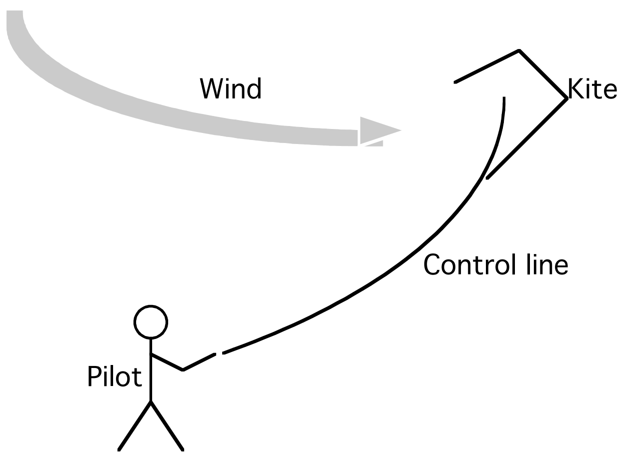

I'll introduce one more component later; it is a device which I call a manipulator, intended to help the pilot to manage a control line composed of several strings, but it adds nothing really new to the system.

The components interact in several ways, many of which are likely to be important in controlling the kite. I think that the table below includes the complete set of significant interactions – or perhaps more than the complete set, but all the interactions listed seem to have some plausible part to play in the overall control system.

From	*To*	*Interface*	*Function*
Pilot	Control line	hands	control
Control line	Pilot	hands	feedback ?
Control line	Kite	attachment	control
Kite	Control line	attachment	feedback ?
Wind	Kite	surface	environment
Kite	Pilot	vision	feedback

The "attachment" is just the joint between the control line and the kite.

I have assumed that the control line is composed of one or more strings, as suggested by the sag shown in the diagram; whatever the natures of the strings, I assume (in the spirit of kites, so far as I understand it) that they can operate only in tension, and can exert no tangential force.

WHAT HAS TO BE CONTROLLED ?

A kite in space has six degrees of freedom, but, because of the nature of the system, not all are independent, and not all are of primary importance. A radius from the pilot is defined fairly precisely by the length of the control line, and relationships between the kite's orientation and its movement – and so, perhaps, an equilibrium position, though I'm not too sure about that – are determined by the wind speed and direction.

Guesswork suggests that it is sufficient to concentrate on two degrees of freedom, regarding the orientation as of secondary importance, and the kite's position on a sphere centred on its pilot as the primary variable. The pilot's control requirements can therefore be summarised as the position of the kite in a two-dimensional space, and the natural way of thinking of these requirements might be expressed in terms of two components :

Vertical coordinate	Real number
Horizontal coordinate	Real number

This, or something like it, is a representation of what the pilot wants to control in terms comprehensible to the pilot; it is what I have called[1] the *human-form* vocabulary, expressing the pilot's requirements for communications in a form which corresponds to his ways of thinking. The essential features of the flow of control information could be described in these terms :

Position (Pilot) →

Signals (Control line) →

Position (Kite)

The objects we would like to control can only rarely understand human-form vocabulary directly, so a general problem is that of translating our wishes into terms acceptable to the objects. The kite is no exception. Its position must be controlled indirectly, as the control line can not directly apply the tangential force we would need to impose our desired coordinates directly. Instead, we must define some sort of *control vocabulary* which the pilot can convey through the control line to determine the kite's motion, and thence its position.

Even if the control line was sufficiently rigid to move the kite sideways, the kite's position remains special in that the kite cannot "understand" its position; there is no state property of the kite itself which is related to position. The kite can "understand" motion, in the sense that it has certain states which cause motion, but it has no states which cause position. Any *requirement of position must therefore be satisfied through visual feedback to the pilot and corresponding corrective action if needed. To emphasise the point, observe that a position-conscious kite could be built; it would be possible in principle to equip a kite with a GPS and perhaps an altimeter, which could then be linked to an automatic guidance system. Then we could transmit the desired coordinates along the control line and the kite would "understand" them well enough to move to the required spot. This would presumably take all the fun out of flying kites.*

More precisely, we must define a set of vocabularies, for the interfaces between pilot and control line, and between control line and kite, are passive, and cannot translate information from one vocabulary to another. The necessary translations must therefore be accomplished by the active system components – in this case, the pilot and the kite. (String is not very intelligent – but the manipulator which I shall introduce as a part of the control line does contribute a little to the translation.)

In what follows, I shall analyse the progress of the control information, paying attention to its encoding at various stages of the control system. The basis of the analysis throughout the description will be that the encoding technique used must be sufficiently expressive to encode all the control requirements, or the system will fail to meet its specification.

The sorts of motion which can be undertaken by a kite are determined by its aerodynamics, of which I know practically nothing. It is clear, though, that there must be two distinct modes of motion, or the exercise of two degrees of freedom would be impossible. In principle, any two linearly independent (?) motions will do; for purposes of discussion, I'll assume an approximation to Cartesian coordinates.

The vocabulary which must be "understood" by the kite is therefore that needed to describe the desired motion of the kite :

Vertical motion :	up	stop	down
Horizontal motion	left	stop	right

The communication system must therefore be able to convey to the kite (at least) two distinct signals which can be interpreted *by the kite* in such a way as to effect satisfactory control over its motion so that all points over some practicable range determined by the wind speed and direction are accessible. The result is positional control over two continuous degrees of freedom.

To convey these instructions to the kite, they must be expressed in terms of the signals which can be transmitted through the control line – the control line's machine-form vocabulary – and these must finally be interpreted by the kite as instructions in its own control vocabulary, shown in the table above.

What can we say about the control line's vocabulary ? At the moment, not much, because we haven't defined the channel at all precisely, but there is one general point which is interesting to discuss at this stage. It is that the control line's fundamental vocabulary is essentially binary. Each string in the line can be in one of just three states, *slack*, *taut*, and *stretched*, where a string is taut in the normal position, with the kite in equilibrium under its weight, the wind, and the control line forces. Further, only one of these – the taut state – is stable, and must in practice be regarded as the null position. Using the other two states, we can impose two sorts of signal on the line. If we relax a taut string so that it becomes slack, the kite will move to a new equilibrium position in which the string is taut again; similarly, if we pull on a taut string to stretch it, the system will again relax to a new equilibrium position. In practice, then, the effect of the string is to transmit the movement of the hand to the kite, and I shall usually regard this movement as the effective control line vocabulary, not worrying a lot about the tensions unless I have to.

The control line is an oddity among communications media in this respect, in that nothing very obvious, even in an abstract sense, passes from one end to the other. There is a notion of direct action in the transmission of sound waves, light waves, radio waves, fingers pushing keys, and electrical transmissions which isn't present with the strings. The reason is that all the kite's motive power comes from the wind; the strings merely impose a constraint on its motion.

HOW CAN THE SIGNAL BE ENCODED ?

The quick answer to that question is "indirectly". I've already pointed out that the control line – the only control channel – cannot convey the tangential forces which I have just defined as the kite's vocabulary. These forces must therefore be derived from some other interaction listed in the table – in this case, obviously the interaction between the kite and the wind. The control line must therefore be employed to do something to the kite which will affect its interaction with the wind. I shall describe this as the control line's changing the kite's *state*.

What sort of state is important ? In principle, any property of the kite which affects the interaction between kite and wind is potentially useful. The immediately obvious property is the orientation of the plane of the kite with respect to the wind :

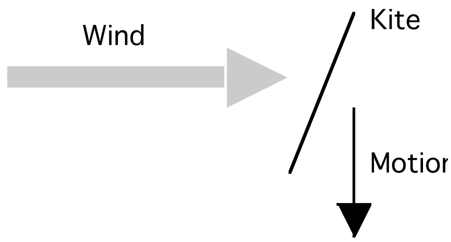

Therefore, by independently controlling the pitch and yaw of the kite with respect to wind direction, the required control can be effected. This is possible with a **three-string** control line, which one could imagine like this :

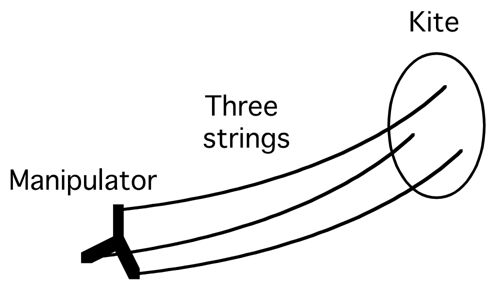

Using this form of control, the pilot must express his control requirements in terms of the kite's orientation, and adjust the manipulator, also by rotations about axes perpendicular to the control line axis, to the required orientation. The manipulator acts as a transducer, which translates the required orientation into string tensions, resulting in changes to the kite's pitch and yaw under the force of the wind until the strings tighten again and the tensions rebalance with the kite in its new orientation. The oblique orientation of the kite with respect to the wind results in a transverse force, causing the kite to move in the required direction; using visual feedback from the kite, the pilot can thereby steer the kite through its accessible space.

(Yes, that's rather obvious, but I've spelt it out in detail to underline the complexity of the process. It's made much simpler than it could be for the pilot by using the manipulator, which automatically takes care of translating the pilot's requirements into tensions, which are the real machine-form vocabulary, which are reflected as relative positions by the process I described above.)

The whole communication process can be described in this way :

Position (Pilot) →
Desired motion (Pilot) →
Kite orientation (Pilot) →
Manipulator orientation (Pilot) →
Manipulator orientation (Control line) →
Tensions (Control line) →
State (Kite) →
Motion (Kite) →
Position (Kite)

This presentation brings out the sequence (or, better, a plausible sequence) of encodings used in the system, and shows where the translations between encodings must be effected.

I've drawn a circular kite to emphasise that with this form of control only the orientation of the kite plane relative to the wind is important; the kite's vocabulary in this system is its pitch and yaw. It's a comparatively complex form of control, though, because you do need the three strings; the change of orientation is controlled by the difference between the string tensions, and to control two orientations you must have two independent differences.

There is another way to get two independent differences. I've assumed so far that the strings forming the control line remain parallel, so that all the tensions operate in parallel, and only their differences are important. What happpens if you twist the manipulator, as if it were a steering wheel ? Now the strings are no longer quite parallel, and we should use vector composition of their tensions. The result will be in general a force (much the same as before) and a torque, not quite the same as before, because it might now include a component about the control line direction which tends to roll the kite about the control line axis. For a circular kite, that makes no difference, but with a less symmetrical body the behaviour might change.

Consider the forces acting on a kite of traditional shape. This is bilaterally symmetrical about its vertical axis, but not symmetrical about its horizontal axis. In a head-on wind, there will therefore be no resultant horizontal force, but there is likely to be a resultant vertical force. (That's less likely if the phenomena concerned are linear; if they are, it doesn't much matter, because you can get an equivalent effect by changing the kite's pitch, but this is a neater argument if it works.)

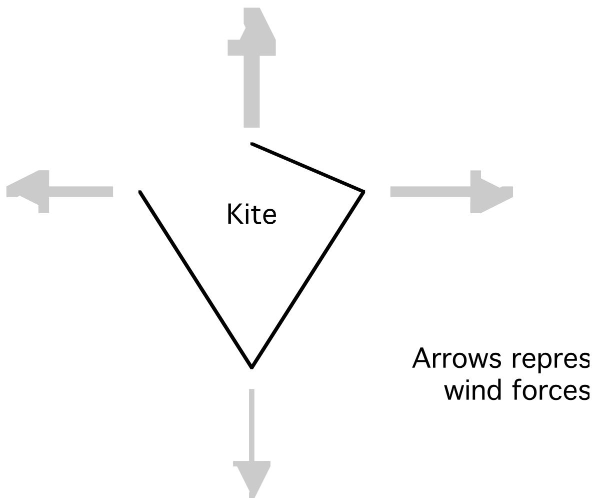

In this orientation, there is a resultant vertical force, but no resultant horizontal force; but if the kite is rotated somewhat in its plane, the resultant force – still along the long axis of the kite – contains a horizontal component.

With this in mind, one could imagine a system involving a **two-string** control line, using the longitudinal tension difference to control one oriention of the kite plane relative to the wind as before, but using the torque around the control line axis to roll the kite and thereby to redirect the resultant force.

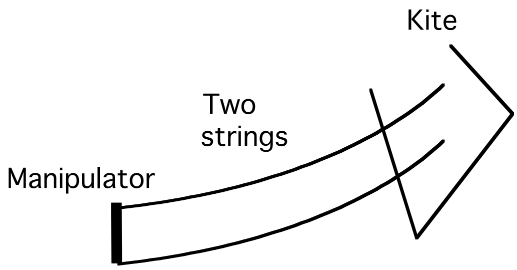

I've drawn the two strings in a vertical plane, to give direct control over the upward component of the wind force. That seems to me to be sensible. (It also guarantees that the result is still valid even if my hopes about the non-linearity of wind forces on a plane surface are dashed.) Rotating the manipulator will now rotate the kite a little, which can cause movement to left and right.

The steps in the communication process are much as in the diagram for the three-string system, but there are differences in the encoding carried out by the pilot. The pilot must now encode his desired motion in terms of an manipulator rotation about a horizontal axis perpendicular to the control line direction, which controls in the main motion in the vertical direction, and a second rotation about the control line axis, which imposes a horizontal migration upon the motion. In this case, the two actions will interact to some extent, for any departure of the kite's long axis from the vertical will both increase the horizontal force component and decrease the vertical component; some compensation might therefore be necessary.

What are the prospects for using a **one-string** control line ? Not obviously very good. It is no longer possible to rely on differential behaviour of different strings, and all control information must be conveyed through the single string tension. As this is essentially limited to the three values slack, taut, and stretched, with the meaning of "taut" determined by the orientation of the kite and the wind, that leaves very little room for manœuvre. That's why the traditional one-string kite goes up to some equilibrium position, all being well, and stays there, also all being well.

If you want to do any better, you have to find a way of changing the kite's state using the single string. There is only one way to do that. First, the signal on the control line must be explicitly encoded in terms of *changes* in tension; and, second, the kite must be so constructed that the changes switch it between two or more stable states, so that they affect its behaviour over an extended period. The stability is new; the states I have discussed so far have been the kite's orientations, maintained by the pilot through the control line, but this maintenance is no longer possible using the single string, so the kite must be able to maintain its own state.

There is an existence proof in my bathroom. An electric switch for the shower is set in the ceiling, and controlled by a dangling string. A tug on the string changes the state of the switch between on and off, and the state, once set, is stable. It works. In terms of the kite, a "change-state" message could be encoded as a brief (guaranteed by the wind) change in tension from taut to stretched, or from taut to slack; if the kite can respond to this message by changing state in some way which changes its aerodynamic behaviour, then the required control is possible in principle. It looks something like this :

Position (Pilot) →
Desired motion (Pilot) →
Kite state (Pilot) →
Control code (Pilot) →
Tensions (Control line) →
State (Kite) →
Motion (Kite) →
Position (Kite)

Now, the control line does less – reasonably enough, because it is less complicated, while the pilot does more, for there is a rather more complicated translation from desired motion to kite state, and the current state must be taken into account. The kite state is also more complex, as it still has the orientation components which it had before as well as the new configurations, but these are not easy to express in this notation.

CONCLUSIONS, OF A SORT.

The first conclusion is that Geoff Crumplin's description (if what I remembered *was* his description) makes sense. It is possible using a single string to switch a kite between different configurations (which I had remembered as "distorting the kite"), and thereby to steer it effectively. I didn't suppose him to be deceiving me in any way, but I couldn't see how it could work; now I can. It remains quite possible that I haven't discovered how his kite really *does* work, but that's another question. I am reasonably confident that, whatever it is, it must rely on some basic mechanism not very different from that which I've suggested.

Another conclusion is that much more complex behaviour could be achieved, because once we can transmit digitally coded messages along the string, they can be made arbitrarily complex. What the kite could do with such messages is again left to the imagination of the reader, and there are several interesting questions connected with the bandwidth of the string, the speed of response of the kite, and the kite's interaction with the wind which would have to be resolved.

SOME COMMENTS.

- I've pointed out that the machine-form vocabulary of the control line is essentially digital in nature. Though the human-form requirements are essentially analogue, which requires that there is some analogue component in the communication, this is provided in practice by varying the time over which a signal is maintained. That, in turn, works because we are, in effect, controlling the rate of approach to the desired position rather than the position itself, so we integrate the control signal over time to attain the required end. I don't think that there's any analogue of this in a more conventional communication system where the intention is to communicate between people in terms of words. If you wanted to achieve that sort of communication with tight strings, you would have to devise some sort of digital code.

(Unless, of course, you use the "bush telephone", where the string tension is modulated :

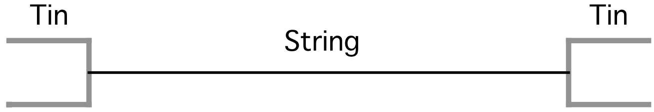

This is quite like the one-string control system, except that the frequencies concerned are much higher. Its limitations are well known, and I'm not advocating it for general use.)

- I haven't discussed radial motion of the kite, but I suspect that it's important : pulling back a kite string quickly, then releasing it gradually, seems to make a difference. (Aerodynamics is notoriously non-linear.) This doesn't materially affect the reasoning I've given, and the essential problem remains two-dimensional, but it does afford an additional channel of control. I just don't know enough about kites or aerodynamics to pursue this notion any further.

- The first two steps (three-string and two-string control) are obvious, even if I've got it wrong, and here serve mainly to introduce the description technique. That's because it's easy to encode the message to be transmitted in the control line's vocabulary. The difficulties arise in the third case, the one-string control, where the communications channel appears to be too narrow to convey the complexity of the control signal required. The systematic approach then focuses attention on the real problem, and leads fairly directly to a plausible solution. In other areas of communication, there is an analogous phenomenon : we don't worry much about the details of communication when we have full command of the conventional media of speech, writing, and drawing; but for people with communications disabilities the problems become acute, and I hope that this technique will be similarly illuminating in such cases.

- I pointed out in a comment that the end result of the control process – the kite's position – was not comprehensible to the kite itself. There is an analogy between this feature of the kite's behaviour and our attempts to communicate using devices such as word processors. In each case, we are attempting to cause inanimate bodies to behave in a way which satisfies some human requirement. Our word processors help us to control what the machinery can do – put dots of ink on paper, or dots of light on a screen, in certain more or less organised ways – but it remains our responsibility to use the tool effectively, because only we are (sometimes) able really to understand what the results is supposed to do and how it works. Similarly, the position of the kite is only significant because someone wants it there; even with my advanced position-control attachment, it's no good telling the kite to go somewhere æsthetically satisfying.

"So what ?", you might reasonably ask. So it's important to make sure that you specify the requirements carefully, I answer. If I am satisfied that being able to control the position of the kite will be sufficient to satisfy my needs, all well and good, but it does require that I learn to fly the kite. Similarly, communicating with people using a word processor requires me to learn both to use the word processor and to use written language effectively. (That's more complicated than a kite, I think.) But if we want more effective computer software, we have to begin with a specification which is more than putting characters on paper.

REFERENCES.

1 : G.A. Creak : *Reaching beyond words in rehabilitation computer systems*, unpublished Working Note AC96 (May, 1996).

2 : G. Crumplin : in conversation, June, 1996.