

Articulating Tangible Interfaces

Alan F. Blackwell

University of Cambridge Computer Laboratory
William Gates Building, 15 JJ Thomson Avenue
Cambridge, UK
Alan.Blackwell@cl.cam.ac.uk

Darren Edge

Microsoft Research Asia
5F Beijing Sigma Center, 49 Zhichun Road
Beijing, China
Darren.Edge@microsoft.com

ABSTRACT

The majority of Tangible User Interfaces (TUIs) consist of rigid objects that are either held in the hands, or arranged relative to each other on a horizontal or vertical surface. In this paper we consider the design space of TUIs that can be created by moving parts of an assembly relative to each other – creating articulated tangible interfaces. An analytic approach to this design space allows us to identify the potential applications and trade-offs of TUIs that include mechanically articulated parts.

Author Keywords

Tangible Interfaces, Mechanisms, Kinematic Pairs

ACM Classification Keywords

H5.2 Interaction Styles

INTRODUCTION

The great majority of Tangible User Interfaces (TUIs) are constructed from solid, rigid objects. These can usually be arranged on a surface, and sometimes stacked on top of one another or placed inside containers. However the individual parts of a TUI seldom incorporate any internal mechanical complexity. In particular, it is rather unusual for TUIs to include articulated joints that act as information-carrying “structural variables” [1]. We believe this is a significant design opportunity – one that will benefit from an analytic framework to understand and generate design concepts.

In this short paper, we present a summary of the cognitive benefits that can arise from the use of articulated joints in tangibles. We then present an analysis of the range of joints that are mechanically possible, with consideration of the trade-offs that are inherent in the different kinds. Finally, we describe a case study design exercise in which we set out to create an articulated TUI as our primary design goal. The analysis draws on previous work that presented a framework for describing trade-offs in TUI design.

COGNITIVE BASIS

In previous work we have described TUIs in terms of research developed from diagrammatic notations and visual

languages, as “manipulable solid diagrams” [2]. The cognitive benefits of diagram use arise from the fact that, unlike symbolic notations [13], the constraints of arranging elements on the surface can be used to reflect constraints of the represented domain. This leads to “free rides” [12] in interpretation and “law encoding” [3] in construction, such that the representation itself removes the need for some mental operations.

Existing TUIs do make use of some straightforward constraint mappings of this kind. Two that are particularly common in diagrams are containment and ordering on the plane. If region A contains region B and B contains C, then we know that A also contains C. Similar readings are possible for “above” and “left-of”. These cognitive efficiency benefits apply directly to TUIs that allow containment of physical objects within each other or arrangement of objects on a surface.

These kinds of cognitive benefit are “natural”, in the sense that they are such a commonplace feature of the perceived world that they are hardly noticeable, while also being processed extremely efficiently by human perceptual systems. As with visual languages and GUIs, TUIs include many diagrammatic aspects that support “natural” exploitation of such mapping constraints.

However, TUIs offer a further opportunity for cognitive efficiency that is not supported by the surface markings of diagrams and GUIs. Unlike surface markings, TUIs can be held in the hands. This can convey additional information, not only through inherent material properties such as weight and texture, but also through familiarity with simple mechanical behaviours such as sliding and rotation. This kind of mechanical intuition, or tacit “knowledge” in the hands, has not yet been investigated by cognitive scientists to the same extent as has been done with visual perception. Nevertheless, we believe that the same analytic approach can be applied to user experiences of mechanical behaviours, as we do in the following section.

MECHANICAL BASIS

We are particularly interested in opportunities to represent continuous digital values by changing the physical configuration of an object. The physical configuration of any mechanism can be decomposed into a relatively small number of fundamental kinematic pairs of elements, as analysed extensively in the 19th century by Reuleaux [11]. The two parts of each kinematic pair are constrained in a

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
TEI 2009, February 16–18, 2009, Cambridge, UK.
Copyright 2009 ACM 978-1-60558-493-5/09/02 ...\$5.00.

way that prevents them from being completely separated, while also allowing them to move relative to each other. The resulting degrees of freedom (DoF) in relative movement between the parts can thus be used to represent one or more continuous digital values.

When considered as a basis for the design of physical-to-digital transducers, the different types of 1 DoF kinematic pair can be seen as different instantiations of the same concept – that of cooperating “threads” that mate together. The prototypical example is a nut and bolt, which together describe a “screw pair” (Figure 1, centre). The relative motion of such a pair describes a continuous helix, corresponding to what we understand as “twisting”. If we imagine the pitch of the screw threads getting increasingly shallower, the nut will eventually turn without moving along the screw, becoming a “revolute pair” corresponding to “turning” motion (Figure 1, left). Alternatively, if we imagine the threads stretching out until they become grooves along the length of the screw, we get a “prismatic pair” where the nut traces a “sliding” path along the screw without turning (Figure 1, right). Hence, screw pairs, revolute pairs and prismatic pairs are the fundamental constructions for mechanically representing one-dimensional values in tangible interfaces, with twisting, turning and sliding being the fundamental actions for manipulating them.




Figure 1. Design space available for one degree-of-freedom kinematic pair, ranging from rotary joint (revolute pair) to slider (prismatic pair).

An illustrative collection of basic kinematic pairs is shown in Figure 2. We do not directly address the question of how digital values can be derived from these pairs, although there are many possibilities. Because kinematic pairs are the basic elements of all mechanisms, the design of electro-mechanical systems has required the invention of transducers that detect movement in all these degrees of freedom (and perhaps also actuate them, as in the Topobo system [10]). A catalogue of industrial sensors offers many possibilities. However, in our own work, we have used camera sensing and simple image processing algorithms to detect the configuration of the parts. Knowledge of the motion constraints, and of the geometry of the rigid parts, enables simple template-matching approaches to inference of the physical joint configuration.

DESIGN OPPORTUNITIES AND TRADE-OFFS

In previous work, we analysed kinematic pairs as one example of how the Tangible Correlates [4] of the Cognitive Dimensions [6] can be used to analyse design trade-offs in TUIs. We summarise those findings here. We can make a comparison between the three 1 DoF kinematic pairs by considering simple expressions of each. A sliding pair can involve one part sliding *within* another – a “position-slider” – or two parts sliding *outside* one another, making a “length-slider” (or telescopic device, such as the stretchable square in the Bricks TUI made by Fitzmaurice, Ishii & Buxton [5]). A turning pair can be a direct rotation device – as in a “knob” – or an indirect rotation device, where rotation of the “joint” is a consequence of the movement of the joined elements. In an analogy to the linear arrangement of sliders, “position-screws” represent some quantity by the position of a nut on a bolt, and “length-screws” by the extension of a bolt out of a threaded cavity (like a swivel chair which uses a screw pair for height adjustment). Other expressions are possible, but they are generally more constrained versions of these six basic types (e.g. a “hinge” is a limited “joint”).




Figure 2. Basic expressions of 1DoF kinematic pairs. Note that while our own prototypes have used optical sensing to detect the configuration of relatively straightforward mechanical assemblies, real applications would require embedded sensors to take advantage of manual interaction skills and provide robust real-time digital feedback.

The most fundamental trade-off within the design space of Figure 2 is between two dimensions: the bulkiness of an articulated tangible (its spatial extent in three dimensions), and its rigidity (its resistance to changes in configuration). The compact linear form of a position-slider means that multiple sliders can be placed side-by-side, for simultaneous operation with simple hand movements. The interface of audio control devices such as graphic equalisers and mixing desks is a good example. In contrast, a length-slider has slightly more bulkiness in one dimension due to the variation in overall size. It also requires two hands to operate, increasing its representational rigidity. A knob can have even less bulkiness than a position slider, but takes slightly more time to operate since there are fewer

tactile cues as to the current value – a marginal increase in rigidity. A joint requires two hands to operate and so has a similar rigidity to a length-slider, but takes up varying amount of area in two dimensions depending on the joint angle, increasing its relative bulkiness. Screw pairs generally have less inherent bulkiness than sliders due to the use of an extra dimension when expressing the single degree of freedom (think of a screw pair as a coiled-up slider). However, they have more inherent rigidity than the other pairs, because many rotations are required to achieve a given translation. Figure 2 shows these trade-offs.

COMPARISON TO EXISTING DESIGN OPTIONS

Based on our comparison of the interaction potential of articulated tangibles with that of traditional TUIs in which separate rigid objects are arranged and rotated on a plane surface, we can tabulate our design guidance (Table 1).

Continuous Attribute Conditions	Recommended Tangible Syntax
Low space constraints One or two attributes to express Tokens unlikely to share the same values	Location on plane
One attribute to express Little movement of tokens necessary	Orientation
Side-by-side attribute comparison Rapid operation	Position Sliders
Peripheral comprehension Rapid operation	Length Sliders
Side-by-side attribute comparison Accurate operation	Position Screws
Peripheral comprehension Accurate operation	Length Screws
Multiple related controls (coaxial or linear) Rapid operation	Knobs
Multiple related controls (coaxial or linear) Peripheral comprehension	Joints

Table 1. Summary of syntax recommended for continuous attributes in different circumstances

There are many other tradeoffs associated with these kinematic pairs, which we can also express using the Tangible Correlates of the Cognitive Dimensions [4]. Joints can be composed into a linkage, providing potential for adaptability. Length-sliders and length-screws have the greatest degree of role expressiveness, as they represent quantity by changes in physical size. Screw pairs are difficult to change accidentally, and so display the least shakiness. Position-sliders and position-screws have a greater degree of juxtaposability when arranged for side-by-side comparison. Pairs in which only one part touches a surface (position-sliders, knobs and some position-screws) also have less rootedness than those in

which both parts rest on the surface – they can more easily be moved without affecting their configuration.

Knobs are a special case, being the only one-degree of freedom kinematic pair that is “stateless” (after a manipulation, the starting position of a knob can be interpreted as the same as it was before that manipulation). Knobs have relatively low rigidity and bulkiness and can also be tailored in different ways. Although a simple knob can only express values within the range of a single revolution (about as expressive as a position-slider), knobs can also be augmented to track the number of revolutions, for example with an array of lights. This simulates the information range of a screw pair. A more abstract virtual layer might allow the knob to exploit its free-turning property, allowing an unbounded range to be represented. However, getting to a value outside the expected range might be a time-consuming if the angular increment of the knob is inappropriate. In this case we might use coupled combinations of knobs, for example controlling logarithmic increments of 1000, 100, 10, and 1. This may result in faster and more accurate control, but the physical state no longer corresponds directly to the value controlled, introducing another system of tradeoffs between rigidity, adaptability and structural correspondence – the similarity between the physical and digital structures. Alternatively, one knob might define a multiplier ratio for the other (like the front and rear gears on a bicycle), as in the SeismoSpin device [9] designed to navigate time on a scale of minutes to decades.

EXPLORATORY DESIGN EXERCISE

As a design exercise to explore the possible interpretation of articulated tangibles, we considered the two rotational degrees of freedom that are provided by the two joints of a three-bar linkage. After an initial exploratory phase investigating the articulated joints used in conventional mechanisms and tools, we constructed a variety of low-fidelity prototypes, including those shown in Figure 3.




Figure 3. Prototypes of articulated tangibles using a three-bar linkage, shown with ruler for scale. Construction techniques: (a) art foam; (b) balsa wood; (c) Perspex with paper overlay

We then built a simple image-processing framework so that we could experiment with interactive possibilities. Our system used an overhead camera setup to track the configurations of multiple such articulated tangibles on the

surface below (Figure 4). Our real-time image processing fits a model of the assembly geometry to the observed scene, allowing the system to continuously track both the positions of tangibles on the surface and the angles of their articulated joints.




Figure 4. System configuration for simple interaction with our prototypes: an overhead webcam is pointed at the desk surface on which articulated tangibles are manipulated.




Figure 5. Alternative layer structures allowing different parts of the articulated object to rest on the table surface & Mechanical alternatives to creating articulated hinges with variable internal friction

We experimented with different physical arrangements, finding that it was possible to create a three-bar linkage that could be placed with one hand, while still allowing independent control of the two free ends (moving them with the index and little fingers respectively). This is not necessarily comfortable to use – bimanual interaction may be easier. However, the physical form of the articulated tangible can be made to facilitate either one-handed operation (by having a centre section that rests stably on the table surface, as in Figure 5a), or two handed operation (where the two free arms rest on the table surface, as in Figure 5b). A further solution (Figure 5c) allows all parts of the tangible to rest on the surface, but increases the likelihood that friction on the surface will modify the joint angle while moving the whole assembly across the table (a trade-off between rootedness, rigidity and shakiness). The prototypes of Figure 5 also demonstrate different approaches to creating rigidity in the joints. Whereas the

joints of Figure 5a and 5b use nut and bolt fastenings to apply pressure parallel to the axis of rotation, Figure 5c demonstrates how internal joint pressure perpendicular to the rotational axis can create a torque force with similar frictional effect. Through our physical prototyping and experimentation, we found that it was useful to include sufficient rigidity in the joints such that the whole object could be placed and rotated on a plane surface while preserving the joint configurations.

Iconic versus Symbolic Articulated Mappings

These articulated tangibles can be used directly to express two continuous attributes in an iconic manner, or indirectly as a means of transforming the object from one symbolic configuration to another (e.g. I-shape, L-shape, C-shape, Z-shape – see figure 6). The intra-object decisions as to how these degrees of freedom should be utilised and interpreted could reasonably be delegated to the user, without affecting usability of the system at the higher, inter-object level.




Figure 6. Possible configurations of a 3-link articulated tangible (left) and a 4-bar linkage (right)

To consider ways in which these linked bar arrangements might be used, with 3 rotary joints, it might be possible to navigate RGB or HSV colour spaces. With 2, we could specify a value with one joint and its uncertainty with the other. An example could be for task time estimates – longer times are likely to be more uncertain, therefore deviations from a symmetric shape would indicate deviations from this norm. Crude graphic equalisers or probability distributions could be made based on the shape of the whole linkage. Hierarchical variables could also be mapped, e.g. using the first joint to specify a day, the second joint to specify a time of that day. Or one joint could specify the ‘gain’ of the other when used as a control, as discussed with respect to knobs. The two could also be used programmatically: if the tangible represents some effect, the first joint could specify the *delay before* the effect, and the second the *duration of* the effect. Recording TV and setting a musical alarm are things that could be managed in such a way.

Overall, one 3-link articulated tangible supports five continuous degrees of freedom – the X and Y location of the tangible on the table, rotation of the whole tangible, and the angular positions of the two independent rotary joints. We continued our design investigation by using the linkage tracking system we had created to provide tangible control in a novel image manipulation application, discussed next.

EXAMPLE APPLICATION

The purpose of this experimental project was to explore the possibilities for mapping 3-link articulated tangibles to a real application. We used two of these tangibles, thereby providing 10 degrees of freedom in total. Our aim was to design an interaction paradigm in which the 10-channel tangibles provided full support for an application, without use of keyboard and mouse. (In a practical application, the keyboard and mouse would probably be used alongside our novel interface, but this self-imposed constraint allowed us to explore a wider range of interaction parameters).

We chose to work in the general area of graphical image manipulation, mainly because our research group includes both interactive graphics and HCI expertise. Whereas earlier work on graphical curve manipulation used a single high DoF curve device [7], our linkages are much more constrained. Nevertheless, there are still many possible ways that the 10 control channels could be mapped to image manipulations. After integrating our simple real-time vision framework with some image processing tools, we rapidly prototyped a variety of alternative applications, settling on image warping.




Figure 7. Trade-off between direct representation and “edge lock” mode for selecting a profile, which combines advantages of rapid yet inaccurate action by users, with slow yet accurate selection by the computer: (a) “approximate” edge profile specified by position and configuration of the white “physical” tangible; (b) “accurate” edge profile calculated as the nearest set of edges matching the approximate edge profile, and drawn as the green ‘digital’ tangible.

In this application, one of the two tangibles is used to map directly to the shape of some feature in the image. The image itself appears on a computer screen. This allows the tangible to be used in indirect “mouse and pointer” style, where the tangibles on the tabletop control a screen representation that is mapped to their shape and position in more ways than simple position correspondence (in contrast to direct tangible design approaches such as video overlay, augmented displays, or rear-projected surfaces). The application supports two modes – one in which the screen representation is mapped directly to the current position and shape of the tangible, and another “edge lock” mode in which the position is adjusted to snap to a salient profile within the image (detected using computer vision methods). This “edge lock” mode is shown in Figure 7.

After one of the tangibles has been moved into position to specify a desired profile within the image, the user can switch the application into warp mode by briefly placing his or her hand over the first tangible. When in warp mode, moving that tangible then causes the image to be distorted in real time in accordance with the modified profile, allowing simple and intuitive transformations of the kind shown in Figure 8.




Figure 8. Images warped in response to linkage manipulation

The second tangible can be then used to adjust two system parameters, in two different interaction modes. When the two arms are placed perpendicular to each other (Figure 9), changing the relative position of the perpendicular arms “twists” the picture either more or less, by changing the radius over which the warp is applied. When the two arms are oriented parallel to each other (Figure 10), pulling them apart causes the on-screen profile of the first tangible to be “stretched”, changing its scale relative to the image.

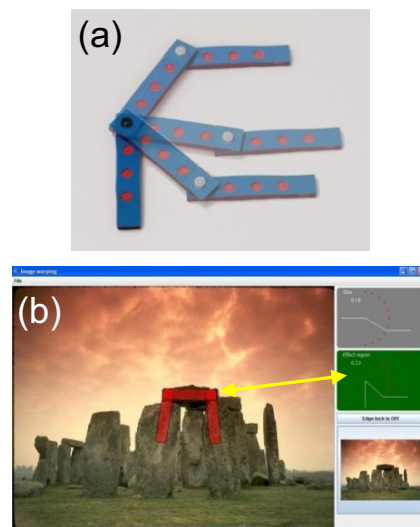


Figure 9. First interaction mode of the second tangible – “twisting” to extend the warp influence parameter: (a) The physical manipulation; (b) the effect of a small twist parameter (compare to large twist parameter of Figure 10c).

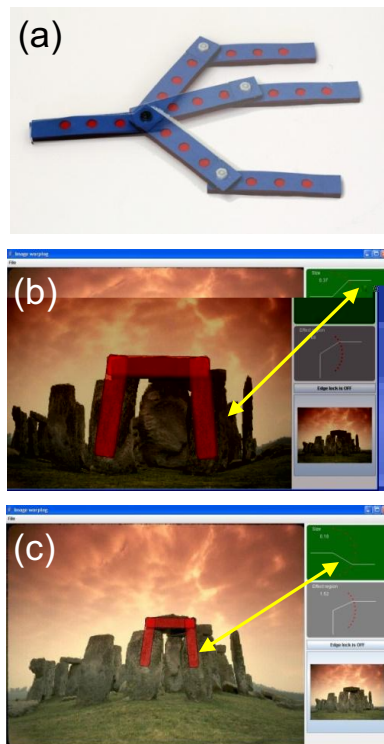


Figure 10. Second interaction mode of the second tangible – “stretching” to change the scale parameter: (a) The physical manipulation; (b) the effect of a large stretch parameter; (c) the effect of a small stretch parameter.

FURTHER DESIGN OPPORTUNITIES

We believe that there are many more opportunities to explore the design space of articulated tangibles. As one example, the construction technique that we used in our exploratory exercise could easily be extended to 4-link assemblies, providing a significant range of symbolic physical geometries as shown in Figure 6.

CONCLUSION

When we consider the evolution from early text interfaces to GUIs, we see that graphical structures required less effort to interpret than symbolic structures because the natural constraints on the plane meant that fewer possible relations could be expressed, thus simplifying both interpretation and choice of actions [12]. These natural interpretive constraints are even more true when making physical arrangements of solid objects in TUIs, because bodily experience of containment-relations, part-whole-relations, linear-relations and attachment-relations appear to underlie almost all abstract meaning [8].

However when we *articulate* the parts of individual objects that still have physical constraints (such as the three alternative ways in which two parts can be articulated to give one logical degree of freedom), we find a rich variety of new abstract relations to which they can be assigned. We therefore propose that physical abstraction in TUIs might best be implemented at the level of articulated objects,

which can then be located in turn within a more directly interpreted frame of spatial relations as in many existing tabletop TUI systems.

ACKNOWLEDGMENTS

This research was supported by The Boeing Company, by Nokia, and by the Engineering and Physical Sciences Research Council undergraduate bursary programme. We are grateful to Ignas Budvytis and Vilius Naudziunas for developing the linkage recognition and image manipulation software described in this paper.

REFERENCES

1. Blackwell, A.F. (2003). Cognitive Dimensions of Tangible Programming Languages. In Proc. of the First Joint Conference of EASE and PPIG 2003, 391–405
2. Blackwell, A.F., Morrison, C. & Edge, D. (2007). A solid diagram metaphor for tangible interaction. CHI'07 workshop on TUIs in context and theory.
3. Cheng, P. C.-H. (1999). Interactive law encoding diagrams for learning and instruction. Learning and Instruction, 9(4), 309-326.
4. Edge, D. and Blackwell, A.F. (2006). Correlates of the cognitive dimensions for tangible user interface. Journal of Visual Languages and Computing, 17(4), 366-394
5. Fitzmaurice, G.W., Ishii, H. and Buxton, W.A.S. (1995). Bricks: laying the foundations for graspable user interfaces. CHI'95, pp. 442-449.
6. Green, T.R.G. & Petre, M. (1996) Usability analysis of visual programming environments: a 'cognitive dimensions' framework. J. Visual Languages and Computing 7, 131-174
7. Grossman, T., Balakrishnan, R. & Singh, K. (2003). An interface for creating and manipulating curves using a high degree-of-freedom curve input device. CHI '03
8. Lakoff, G. (1993). The contemporary theory of metaphor. In A. Ortony (Ed.), Metaphor and Thought, 2nd edition. Cambridge: Cambridge University Press
9. McKelvin, M., Nestande, R., Valdez, L., Yee, K., Back, M. & Harrison, S. (2003). SeismoSpin: a physical instrument for digital data. CHI'04, 832–833
10. Raffle, H. S., Parkes, A. J., & Ishii, H. (2004). Topobo: a constructive assembly system with kinetic memory. CHI'04, 647–654
11. Reuleaux, F. (1876). The Kinematics of Machinery: Outline of a Theory of Machines. Reprinted by Dover Publications in 1963
12. Shimojima, A. (1996). Operational constraints in diagrammatic reasoning. In G. Allwein & J. Barwise (Eds), Logical reasoning with diagrams. Oxford University Press, 27-48
13. Stenning, K. & Oberlander, J. (1995). A cognitive theory of graphical and linguistic reasoning: logic and implementation. Cognitive Science 19:97-140