Spatial Reasoning Using Sinusoidal Oscillations

Ian Pratt
Department of Computer Science
Victoria University of Manchester
Manchester, M13 9PL
England

Abstract

This paper outlines some preliminary results concerning the use of sinusoidal oscillations to represent vectors in two-dimensional space. The proposed representation scheme permits efficient implementation of translation and rotation, and immediate detection of such relations as collinearity and proximity of points. This scheme is then extended so that arbitrary convex regions of the plane are represented using a pair of signals varying over time. Finally, the advantages of representing convex regions in this way are shown to derive from the resulting ease with which such regions can be translated and rotated in the plane, and—more strikingly—from the simplicity of determining whether two such regions overlap.

1 Introduction

This paper presents an approach to spatial representation such that that: (i) translations and rotations of points and regions can be readily performed; (ii) relations such as collinearity and proximity of points can be efficiently detected; and (iii) the question of whether regions overlap, and, if so, by how much, can be settled with a minimum of computation. For reasons which will become evident, we concentrate on the problem of representing *convex* regions of the *plane*.

The proposed representation system is based on a well-known correspondence between sinusoidal oscillations with a fixed frequency ω, and vectors in two-dimensional space. Section 2 describes how we can represent a vector (equivalently, a point in space), by having a device which emits an appropriate sinusoidally varying signal. (Here, "device" may be either a physical device or a virtual device—i.e. a programming construct.) Section 3 extends these results to show how we can represent any convex region of the plane using a device emitting a pair of signals varying with some fixed time period. Such regions, thus represented, can be efficiently translated and rotated; and, as we see in section 4, given two such regions, it is a straightforward matter to determine whether they overlap. As such, I claim, the proposed system promises to combine the expressive power of numerical co-ordinate systems and their variants, e.g. [2], [3], with the computational tractability of qualitative, object-based spatial representation schemes such as [1].

2 Representing vectors

There is a natural correspondence between the set of sinusoidal oscillations with some fixed frequency ω and the set of vectors in two-dimensional space. Given a sine wave

$$w(t) = a\sin(\omega t + \phi)$$

with frequency ω, amplitude a and phase-lead ϕ, we can associate the vector v whose polar co-ordinates are (a, ϕ).

The key feature of this correspondence is that, given two vectors, the *superposition* of their corresponding sinusoidal waves yields a third sinusoidal wave, also of frequency ω, which corresponds to the vector sum of the original two vectors. More formally:

Theorem 1: Let $w_1(t) = a_1\sin(\omega t + \phi_1)$, $w_2(t) = a_2\sin(\omega t + \phi_2)$ be sinusoidal waves of frequency ω. Let v_1, v_2 be the corresponding vectors (i.e. having polar co-ordinates (a_1, ϕ_1), (a_2, ϕ_2) respectively). If $w_3(t)$ is the pointwise sum $w_1(t) + w_2(t)$, and if v_3 is the vector sum $v_1 + v_2$, then $w_3(t)$ is also a sinusoidal wave of frequency ω, $w_3(t) = a_3\sin(\omega t + \phi_3)$, where (a_3, ϕ_3) are the polar coordinates of the vector v_3.

Textbooks on the theory of AC circuits use this result to calculate the effects of combining sinusoidally varying currents and voltages by adding the corresponding vectors. Here, however, rather than using vectors to reason about sine waves, I propose that we use sine waves to reason about vectors.

If we have a device (physical or virtual) continuously emitting a signal varying sinusoidally over time with frequency ω, that device can represent a vector. The amplitude, a, of the wave represents the length of the vector, and the phase lead, ϕ (measured relative to some internal clock), represents its orientation. The varying signal may be a sequence of numbers, or a continuously varying potential difference, or yet other things besides. In this paper, I leave the details of the implementation open, assuming only that suitable sine waves can be readily produced, superposed, and compared in simple ways.

Now suppose we have two such devices running in parallel and emitting sine waves of frequency ω, representing, respectively, vectors v_1 and v_2. By piping those signals to a unit which superposes (i.e., pointwise adds) them—call it an *addition unit*—we have, in effect, computed the vector sum $v_1 + v_2$ (fig.1).

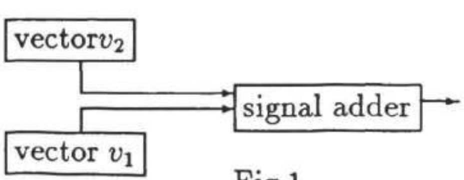

Fig.1

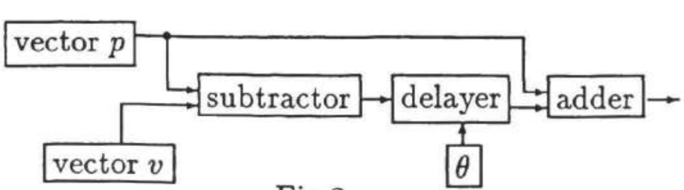

Fig.2

Similarly, *mutatis mutandis*, for vector subtraction: by piping the two sinusoidal waves to a *subtraction unit* which pointwise subtracts one from the other, we have computed $v_1 - v_2$.

To rotate a vector v, thus represented, clockwise through some angle θ (about the origin), we need only increase, by θ, the phase-lead of the wave used to represent v. This corresponds to *delaying* v's output signal by the amount of time $(2\pi - \theta)\omega^{-1}$ (ignoring the first oscillation), again, an operation which I shall take it can be simply performed. To rotate a vector about a point p (other than the origin), it suffices to find $v - p$ (obtainable by piping v and p into a vector subtraction unit), rotate the result through θ about the origin (as just described) and add p (by piping the result together with p to a vector addition unit). The process is illustrated in fig.2.

It is also easy to test for such relations as sameness of orientation, collinearity and proximity. Two vectors have the same orientation if and only if their corresponding sine waves are in phase; so a device able to check that the maxima of two waves coincide could perform this test. To test for the collinearity of v_1, v_2 and v_3, it suffices to compute both $v_1 - v_3$ and $v_2 - v_3$ and then to determine whether these two vectors have the same or opposite orientation. Thus, the arrangement depicted in fig.3 would suffice for such an operation.

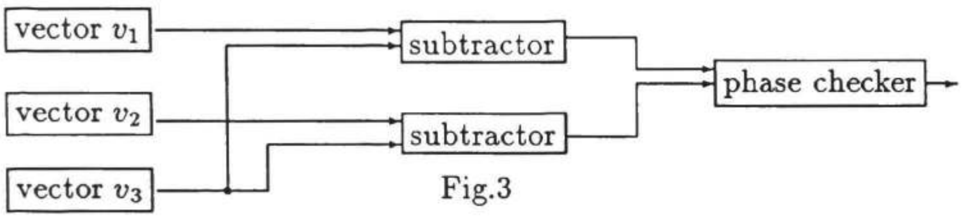

Fig.3

Similarly, two points, corresponding to vectors v_1 and v_2, will be *near* each other if and insofar as the length of $v_1 - v_2$ is small. Therefore, detecting the proximity of points just involves piping the corresponding signals to a subtraction unit and then piping the result to a device which registers the maximum amplitude of an incoming signal.

3 Representing convex regions

Let us begin with a simple example. Consider the line AB between two points (vectors) v_A, v_B, as shown in fig.4 (the sine-wave representation of v_A, v_B is also given).

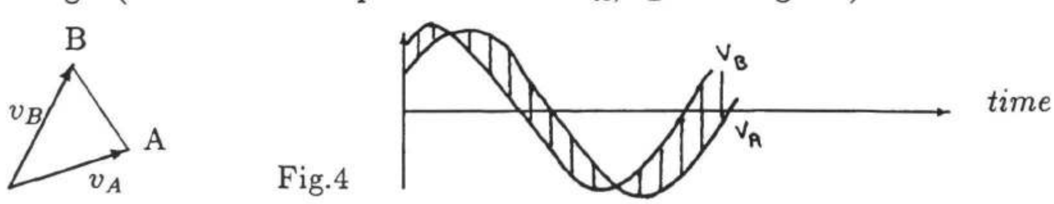

Fig.4

The set of vectors lying on the line AB is the set $\{\kappa v_A + \lambda v_B | \kappa, \lambda \in [0,1], \kappa + \lambda = 1\}$. Now there is a correspondence between this set and the shaded area of fig.4: let $w(t)$ be any sine wave (of frequency ω) and let v be its corresponding vector; then $w(t)$ lies within the shaded region of figure 9 if and only if v lies on the line AB. The task of this section is to generalise this result to show how more complicated regions can be represented.

Definition: A set S of vectors is said to be *convex* iff, $\forall x, y \in S,\ \forall \kappa, \lambda \in [0,1]$ s.t. $\kappa + \lambda = 1$, we have $\kappa x + \lambda y \in S$.

Fig.5 shows some more convex sets of vectors (the shaded regions, *not* just their boundaries) accompanied by the sinusoidal representations of some of the key points on their peripheries. (Points in the plane are identified with the corresponding vectors in the obvious way.)

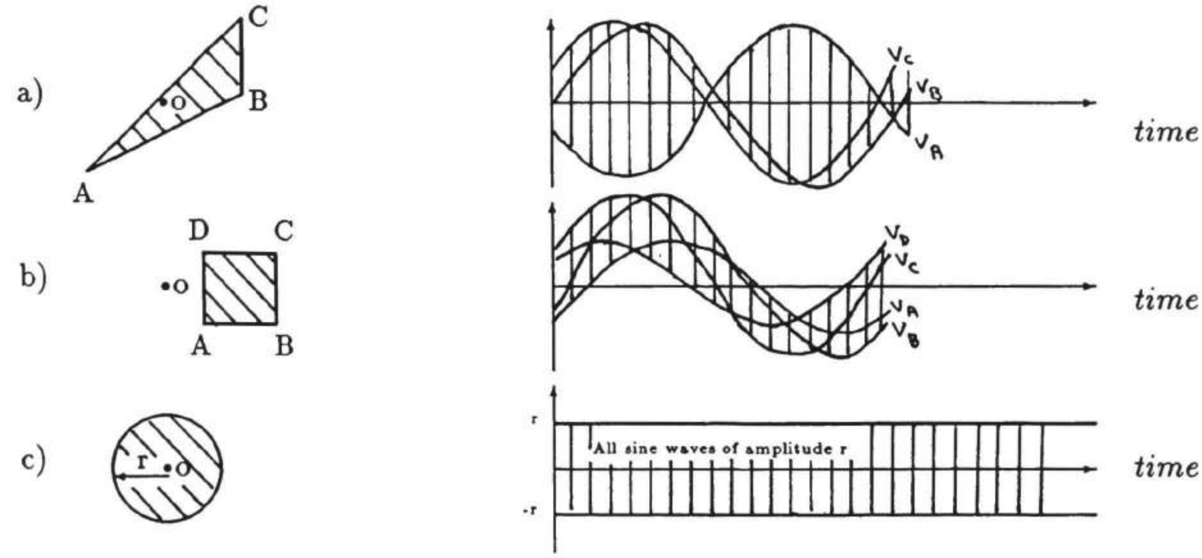

Fig.5 ("o" indicates the origin)

We shall presently show that the shaded regions in the two sets of diagrams correspond, in a sense we now proceed to define.

Definition: If V is a set of vectors, define the *convex hull* of V, to be the smallest convex set containing V.

Definition: If W is a set of sine waves $w(t)$ of frequency ω, define the *sinusoidal hull* of W, to be the pair of functions: $\Sigma(t) = \langle a(t), b(t) \rangle$, where $a(t)$, $b(t)$ are functions defined by: $a(t) = \sup_{w \in W} w(t) \qquad b(t) = \inf_{w \in W} w(t)$.

$a(t)$, $b(t)$ represent the upper and lower bounds, respectively, of all the sine waves in W; if these upper or lower bounds do not exist, we set $a(t) = \infty$, $b(t) = -\infty$, as appropriate.

Definition: If $\Sigma(t) = \langle a(t), b(t) \rangle$, is a sinusoidal hull, and $w(t)$ is a sine wave, we write $w \preceq \Sigma$ (read *w is contained in* Σ) iff $\forall t,\ a(t) \geq w(t) \geq b(t)$.

We can now state the main theorem of this section:

Theorem 2: Let V be a closed set of vectors, and let W be the set of sine waves corresponding to the vectors in V. Let S be the convex hull of V. Let Σ be the sinusoidal hull of W. If v is any vector, and $w(t)$ its corresponding sine wave, then $v \in S$ iff $w \preceq \Sigma$.

In other words, the sinusoidal hull of a set (finite or infinite) of sine waves corresponds to the convex hull of the set of vectors those sine waves represent.

Hence, in order to represent the convex hull of a finite set of vectors, all you need to do is to take the sine waves corresponding to those vectors and pipe them to two devices, one outputting the (pointwise) maximum of all its incoming signals, the other the minimum. (Thus these signals correspond to the functions $\langle a(t),\ b(t)\rangle$ respectively of some sinusoidal hull Σ.) Fig.5a shows roughly what those outputs would have to be in order to represent a triangle. Fig 5b does the same for a square. Such a pair of signals can be used to determine whether any given vector v, represented by sine wave $w(t)$, is in the region in question simply by checking whether $w \preceq \Sigma$ (i.e, whether $a(t) \geq w(t) \geq b(t)$ holds over an entire oscillation).

Not that a sinusoidal hull *need* be generated from a finite number of sine waves in this way: fig.5c shows how the sinusoidal hull corresponding to a circular region centred at the origin is easily generated — it is just a pair of constant signals. But it cannot be generated from any finite number of sine waves. Similarly, for other convex, curvilinear shapes.

Furthermore, *any* device emitting a pair of periodic signals $a(t)$ and $b(t)$, with period $2\pi\omega^{-1}$, can be taken as representing a (possibly empty) convex set. To see this, note that there will either be a maximal sinusoidal hull that lies between these two signals, or else no such sinusoidal hull at all, where a *maximal* sinusoidal hull is one for which the set of sine waves it contains is maximal. And since that maximal sinusoidal hull (if there is one) will pick out some convex set or other, then the original signals, $a(t)$ and $b(t)$, can themselves represent that convex set. As before, to test whether a vector is in that set, we take its corresponding sine wave $w(t)$ and check whether $a(t) \geq w(t) \geq b(t)$ holds over an entire oscillation.

Suppose we want to translate a convex set C by a vector v, where C is represented by the signals $\langle a(t),\ b(t)\rangle$ and v by the signal $w(t)$. Then the translated convex set will be represented by $\langle a(t) + w(t),\ b(t) + w(t)\rangle$. Hence, the translation of convex regions, like that of single points, can easily be effected by means of piping the appropriate signals to an addition unit as described above. Similarly, the rotation clockwise through an angle θ of a convex region about the origin can be effected by delaying the signals $a(t)$ and $b(t)$ by an amount of time $(2\pi - \theta)\omega^{-1}$.

4 Detection of Overlaps of Convex Regions

To detect overlaps, the following definitions will prove expedient:

Definition: Suppose S_1, S_2 are sets of vectors. Define their *inner difference*, $S_1 \ominus S_2$ to be the set $\{(x - y) | x \in S_1, y \in S_2\}$.

Definition: Suppose $\Sigma_1(t) = \langle a_1(t), b_1(t)\rangle$ and $\Sigma_2(t) = \langle a_2(t), b_2(t)\rangle$ are sinusoidal hulls. Define their *inner difference*, $\Sigma_1 \ominus \Sigma_2$ to be the pair of functions $\langle a_1(t) - b_2(t), b_1(t) - a_2(t)\rangle$.

If S_1, S_2 are convex, so is $S_1 \ominus S_2$. Also, $S_1 \cap S_2 \neq \emptyset$ iff $0 \in S_1 \ominus S_2$. Hence, to determine whether S_1 and S_2 overlap, it suffices to determine whether $0 \in S_1 \ominus S_2$. If Σ_1, Σ_2 are sinusoidal hulls, so is $\Sigma_1 - \Sigma_2$. It is formed by subtracting the lower bound of the second hull from the upper bound of the first, and the upper bound of the second hull from the lower bound of the first. Inner differences of sinusoidals hulls can therefore be computed using subtraction units of the kind encountered above. Also, inner differences of sinusoidal hulls correspond to inner differences of convex sets:

Theorem 3: If S_1 and S_2 are convex regions corresponding to sinusoidal hulls Σ_1, Σ_2 respectively, then $S_1 \ominus S_2$ is a convex region corresponding to the sinusoidal hull $\Sigma_1 \ominus \Sigma_2$.

Whence:

Theorem 4: Let w_0 denote the zero-amplitude sine-wave: $w_0(t) = 0$ for all t. If S_1 and S_2 are convex regions corresponding to sinusoidal hulls Σ_1, Σ_2 respectively, then $S_1 \cap S_2 \neq \emptyset$ iff $w_0 \preceq \Sigma_1 \ominus \Sigma_2$.

Since, then, we have a method for computing the inner difference of two sinusoidal hulls, and a means of determining whether a sinusoidal hull contains the zero wave, we can immediately detect whether two convex regions overlap by piping their outputs to an appropriate arrangement of devices.

A caveat. Theorem 4 should be applied with care. Recall from section 3, that *any* pair of periodic signals (with period $2\pi\omega^{-1}$) can represent a (possibly empty) convex set. Given two such pairs of signals, $\langle a_1(t), b_1(t)\rangle$ and $\langle a_2(t), b_2(t)\rangle$, representing S_1, and S_2 respectively, we can define their inner difference in the same way as for sinusoidal hulls. But it is in general *false* that the resulting pair of signals $\langle a_1(t) - b_2(t), b_1(t) - a_2(t)\rangle$ will contain w_0 only if $S_1 \cap S_2 \neq \emptyset$. Theorem 4 requires that $\langle a_1(t), b_1(t)\rangle$ and $\langle a_2(t), b_2(t)\rangle$, be sinusoidal hulls, as defined in section 3.

5 Discussion

On the present view, a complex two-dimensional scene is represented by a collection of "devices", each broadcasting either a single sine wave (representing a single point) or a pair of periodically varying signals (representing a convex region). By piping the outputs of these devices through suitable transforming units, geometrical operations like translation and rotation can be performed. Moreover, by piping the same outputs through yet other devices, tests for relations like collinearity and proximity of points and overlap of regions can be made. The computation involved in these operations and tests is trivial, in marked contrast to the nightmarish problem of determining intercepts of curves expressed as equations in, say, a Cartesian co-ordinate system. It is as if the sinusoidal representation "flattens out" the problem of determining the membership of a point in a convex region, yielding the easier problem of determining whether one signal lies between two other signals over a period of time.

It will be objected that this assessment ignores the difficult task of converting back and forth between a Cartesian co-ordinate representation and the sinusoidal hull representation.

But this would be to miss the point of the present proposal. That proposal is not that regions be represented by equations of bounding curves, and that those equations be converted into sinusoidal hull signals in order to perform transformations and tests on them. Rather, it is that the equations of bounding curves can be dispensed with altogether: all we have is the collection of signal-emitting devices. Thus, to know what a square is, or a triangle or a circle or an oval, our system would have to be able to activate a device outputting the appropriate signals, not state an equation in Cartesian co-ordinates (of course: the device in question might have to store its signals as a pair of equations $a(t) = ..., b(t) = ...$). To repeat: the policy is not to augment traditional co-ordinate representations, but to dispense with them.

So conceived, the fate of the proposed representational system depends on how well it can interface with visual input and motor ouput: no good doing away with co-ordinate representations in favour of sine waves, if it then becomes impossible for, say, a robot to determine the shapes of things it is looking at, or to adjust its movements in conformity with its beliefs about the spatial arrangement of its environment. These are therefore important questions for future work in this area. Other pending questions concern the extension of the proposed system to deal with non-convex shapes, 3-dimensional space, uncertainty and moving scenes. Finally, of course, many implementation details, some non-trivial, remain outstanding.

6 Conclusion

This paper has presented a system for detailed representation of convex plane figures. Its utility derives from the facility it provides to rotate and translate such regions, and to detect overlaps between them. Work remains to be done to extend and implement the system as proposed here. However, the results presented above suggest that, for detailed spatial representations, traditional co-ordinate systems may be considerably less efficient than the current approach.

References

[1] Brooks, R.A: "Symbolic Reasoning among 3-D Models and 2-D images", *Artificial Intelligence* **17**, pp.285-348 (1981).

[2] Klinger, A. and M.L.Rhodes: "Organization and Access of Image Data by Areas", *IEEE Trans. Patt. Anal. Mach. Intell*, PAMI-1 pp.50–60 (1979).

[3] McDermott, D.V. and E.Davis: "Planning Routes through Uncertain Territory", *Artificial Intelligence* **22**, pp.107–156 (1984).