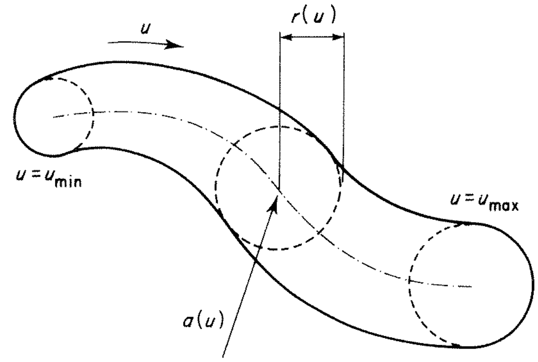

Fig. 10. A swept sphere of varying radius (from Van Wijk).

If $f(x, y, z, t)$ is algebraic, this can always be done using elimination. A similar theory also exists for parametric surfaces.

Van Wijk has applied this method to ray tracing a sphere moving along a cubic space curve with cubically varying radius (see *Figure 10*). For this shape, f is

$$(x - a_x(u))^2 + (y - a_y(u))^2 + (z - a_z(u))^2 - r^2(u) = 0$$

where $(a_x(u), a_y(u), a_z(u))$ is the center of the sphere and $r(u)$ is the radius. In the general case, the resulting implicit equation will have degree $= 2(2d - 1)$ where d is the degree of the trajectory. Thus, for a cubic trajectory the result will be a tenth-degree equation.

Generalized cylinders

A *generalized cylinder* is the surface defined by sweeping a two-dimensional contour along a three-dimensional trajectory. To make this definition exact the position and orientation of the contour relative to the trajectory must also be specified. The most natural method to do this is to orient the contour relative to the Frenet frame of the trajectory. Methods for performing the ray intersection calculation for this primitive are described in [13].

Constructive solid geometry—CSG

One of the most popular methods to model solid objects or volumes is *constructive solid geometry*. Solids of more complicated objects are constructed from simpler objects by performing set operations, usually union or difference. A composite object can be represented as a binary tree where each node contains a set operation and two child nodes which may also be

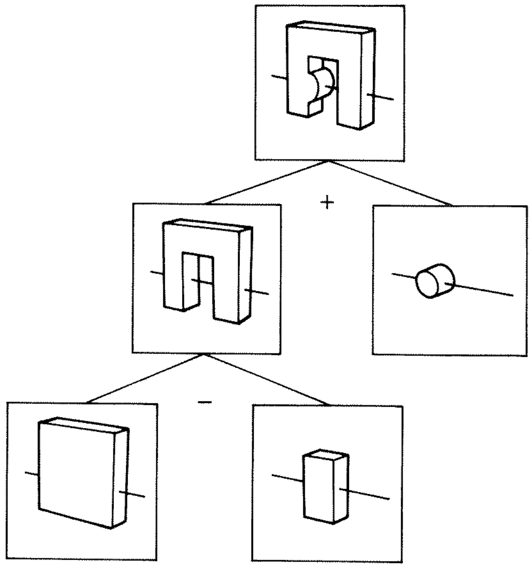

Fig. 11. Example of a solid object formed using CSG (from Roth).

composite solids. The leaves of this tree are primitive objects such as spheres, planes, etc. This is shown in *Figure 11*. One of the advantages of this representational scheme is its completeness; another advantage is the ease with which it can be ray traced.

Evaluating set operations usually reduces to classification of geometry with respect to other geometry [55]. A curve, surface or volume is deemed to be inside, outside or on another solid. When ray tracing, it is necessary to classify a ray with respect to the solid the ray is being intersected against. This intersection path may be represented with a *Roth diagram* (see *Figure 12*) [57]. This diagram is a line from $-\infty$ to ∞ where each point represents a value of the ray parameter t. Each point on this line is either inside, on or outside the solid. The transitions from inside to outside, and vice versa, will be at intersections of the ray with the primitive volumes comprising the solid. If the solid is a primitive, it is easy to generate the Roth diagram. Each intersection of the ray with the solid is shown on the line. Ignoring singularities, any point of the line can be classified by testing an interior point (between two intersections) with respect to the primitive. Sometimes this is unnecessary since it is known *a priori* what the classification of the origin of the ray is with respect to the primitive (for example, if the ray starts at the eye and it is known that eye point is outside all the objects, or if the origin of the ray corresponds to a previous intersection which has already been classified). Once one

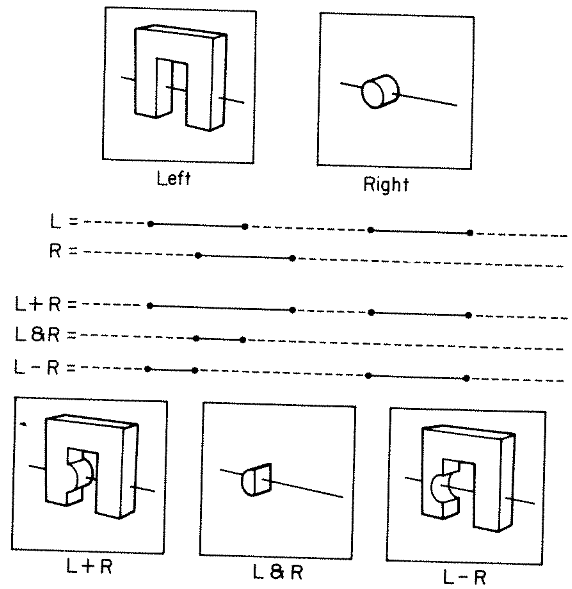

Fig. 12. Example of a Roth diagram for analyzing CSG operations (from Roth).

interval has been classified, all other intervals are alternately classified inside or outside (assuming the surfaces of the primitives are not self-intersecting and there are no singularities). (For a discussion of how to handle the more general problem of curve–solid classification see [68] or [55].) If the surface is orientable, whether a transition is from inside to outside or vice versa can be determined directly from the surface normal at the point of intersection. Note that in order to create a Roth diagram for a solid it is necessary to find all the intersections of a ray with a primitive volume.

Roth describes a method to combine these diagrams using set operations. This allows the Roth diagram for a composite solid to be recursively computed. First the Roth diagrams of the two solids being combined are computed and then these diagrams are merged to form the diagram for the composite solid. The diagrams of two solids can be combined by using the one-dimensional set operation on the line intervals comprising the diagram. This is done in a three-step process. (1) The intersection points from the two diagrams are merged together into a single diagram (this takes time proportional to the total number of intersections). (2) The intersection points are classified depending on the classification of the original two diagrams and the set operation being performed. (3) The interval is simplified by removing intersections which do not result in a change in the classification. The rules for combining classifications are given in *Table 1* based on well-known rules of

Table 1. Rules for combining classifications.

Set operator	L	R	Composite
Union	IN	IN	IN
	IN	OUT	IN
	OUT	IN	IN
	OUT	OUT	OUT
Intersection	IN	IN	IN
	IN	OUT	OUT
	OUT	IN	OUT
	OUT	OUT	OUT
Difference	IN	IN	OUT
	IN	OUT	IN
	OUT	IN	OUT
	OUT	OUT	OUT

boolean algebra. Once again, these rules for combining intervals are only valid when there are no singularities—that is, no ON classifications. If this occurs a table does not suffice and the classification must make use of neighborhood information [68]. Fortunately, when ray tracing this is seldom a problem.

As will be mentioned below, ray tracing is significantly speeded up if the scene is arranged in a tree of bounding volumes. A CSG tree can serve as this tree (although it may not be the best choice) by combining bounding volumes of the subtrees according to the set operation. Usually, all the bounding volumes in the tree are either spheres or boxes so it is reasonably easy to form a new bounding volume from them. The union of the two bounding volumes can be formed by enclosing both of them; the intersection by enclosing the intersection of the two bounding volumes. If the operation is difference, the bounding volume of the composite is the same as the bounding volume of the object not being subtracted.

Roth also observed that the result of a set operation is sometimes known after the ray has been classified with respect to only one of the two child solids. For example, if the operator is intersection and one solid has no intersections with the ray, there is no need to intersect the ray against the other solid. A similar situation occurs when doing a difference; if the ray is outside the positive solid no intersection can possibly occur.

Hierarchical bounding volumes

Many ray tracing algorithms use a hierarchical tree of bounding volumes to either speed up the search for an intersection or to control the generation of