Presented at SPIE Visual Communications and Image Processing '91: Image Processing Conference 1606, pp. 320-334, Nov. 11-13, 1991, Boston MA.

CONNECTIVITY-PRESERVING MORPHOLOGICAL IMAGE TRANSFORMATIONS

Dan S. Bloomberg

Xerox Corporation
Palo Alto Research Center
3333 Coyote Hill Road, Palo Alto, CA 94304

Abstract

Methods for thinning connected components of an image differ in the size of support, type of connectivity preserved, degrees of parallelism and pipelining, and smoothness and fidelity to structure of the results. A unifying framework is presented, using image morphology, of all 4- and 8-connectivity-preserving (CP) transformations that use a 3x3 basis of support on binary images discretized on a square lattice. Two types of atomic CP transformations are defined: *weak* CP neither breaks nor joins components and *strong* CP additionally preserves the number of connected components. It is shown that out of thousands of possible 3x3 hit-miss structuring elements (SEs), in their most general form there are only four SEs (and their rotational isomorphs), for each of the two sets (4- and 8-connectivity), that satisfy strong CP for atomic operations. Simple symmetry properties exist between elements of each set, and duality relations exist between these sets of SEs under reversal of foreground/background and thinning/thickening operations. The atomic morphological operations, that use one SE, are intrinsically parallel and translationally invariant, and the best thinned skeletons are produced by sequences of operations that use multiple SEs in parallel. A subset of SEs that preserve *both* 4- and 8-connectivity have a high degree of symmetry, can be used in the most parallel fashion without breaking connectivity, and produce very smooth skeletons. For thickening operations, foreground components either self-limit on convex hulls or expand indefinitely. The self-limited convex hulls are formed either by horizontal and vertical lines, or by lines of slope ± 1. Four types of boundary contours can result for thickening operations that expand indefinitely. Thickened text images result in a variety of typographically interesting forms.

Keywords: image processing, thinning, thickening, skeletonization, morphology, mathematical morphology, connected components, image connectivity

1 Introduction

The *connectivity-preserving* (CP) image transformations that underlie both thinning and thickening are identical. Thinning algorithms on binary images have a long history in image processing, because of their value in deriving higher representations (and compressed encodings) of the information in a bitmap. Thinned images possess a subset of the original information, that is useful for applications such as segmentation, feature extraction, vectorization, and pattern identification. Thickened images have been of less interest; they allow generation of connected component convex hulls.

There is great diversity both in the methods that have been used to thin image components, and in the results obtained. Most proposed binary thinning algorithms operate directly on the image. Approaches that have been used include: (1) sequential, data-dependent operations (either on the image or on a line adjacency graph representation) acting on component boundaries; (2) sweep/label operations (typically 2-pass) on the entire image, with subsequent operations depending on the type of skeleton to be produced; (3) pipeline on sequential pixels (typically with hardware assist); (4) SIMD parallel on all pixels (with or without hardware assist). In spite of this diversity, at the heart of every algorithm is a procedure that preserves connectivity.

A vast literature on thinning algorithms has accumulated during the past quarter century, and it is amusing to observe that many of the recent publications have the word "fast" in the title, as if without this assurance the reader might assume the proposed method is slow! This proliferation of work on thinning may seem surprising: are operations that preserve image connectivity so complicated that there exist a multiplicity of useful approaches? The answer to this question is in the affirmative, and the purpose of this paper is to provide a simple framework for examining the complexity.

A few comments on various approaches may be useful. The efficiency of a thinning procedure is not simply an intrinsic property of the algorithm; it also depends on both the hardware and the data within the image. Parallel processing can be obtained either through pipelining, where each cycle a new pixel enters the pipeline, or SIMD ("single instruction, multiple data") processing, where each cycle a set of pixels undergo the same operation, or both. Operations within a pipeline are typically local: a pixel can communicate only with a set of its neighbors. The efficiency of a pipeline architecture is proportional to the pipeline depth. Unfortunately, pipeline depth is limited in connectivity-preserving operations using 3x3 local rules, because as the image is transformed by thinning (for example) from each of four directions, each successive operation usually must be applied to the transformed image. This is often referred to as *sequential* thinning. In special cases the pipeline can be extended to the full set of local rules, at an added computational cost of restoring pixels that should not have been altered within the pipeline (Chin et al.[3]). Alternatively, an algorithm with greater parallelism can be constructed by expanding the region of support to 5x5. The pipeline depth can then be increased to a full iteration cycle of all four thinning directions, but at great increase in complexity (Rosenfeld[11]). For sparse images, a pipelined algorithm gains efficiency, relative to SIMD, because only a few pixels in the image need to be processed. In general, SIMD processing is better suited to iterative, sequential thinning and thickening, because the transformed image is always available at the next cycle. The algorithms in this paper are designed to use only logical raster operations, and can thus be implemented either on a very simple SIMD machine or in word-parallel on a general purpose computer. The number of iterations required is proportional to the "thickness" of the largest connected component.

Many of the recent thinning algorithms are intended to be implemented on selected pixels using integer arithmetic. In 1984, Zhang and Suen[19] proposed a method for 8-connected thinning, based on local operations with a 3x3 support. A number of refinements of this method then appeared[18, 8, 7, 6]. These methods differ from each other to some extent in (1) the degree of erosion of free ends, (2) the number of operations required for each iteration, and (3) the size of the support for the local operations. For example, the Zhang and Suen algorithm used a support of 3, but some of the subsequent algorithms implicitly used a support of 4.

Binary morphological approaches to thinning were first described for hexagonal lattices by Golay[4], and more recently summarized by Serra[15]. Maragos and Schafer[9] extended this work in 1986, demonstrating computation of Blum's[2] medial axis skeleton on a square lattice grid.

Stefanelli and Rosenfeld[16], Rosenfeld[12] and Arcelli et al.[1] took an approach to thinning quite similar to the one presented here. In particular, the 1975 papers on parallel thinning algorithms provide insights into both the conditions under which pixel removal can be determined locally, and the constraints on parallelism that must be imposed to preserve connectivity. Rosenfeld[13, 14] describes a particular sequential thinning rule, using 3x3 support, that can be applied to either 4- or 8-connectivity: successively from each side, remove all border pixels that are connected to exactly one connected component that is not an *end point* (i.e., that has more than one pixel within the 3x3 window). [In effect, we are providing a systematic method for constructing parallel boolean implementations of Rosenfeld's rule.] The section on thinning in Rosenfeld's book[14] is also recommended as an introduction. Vincent[17] has recently given an excellent review of skeleton types, along with an efficient sweep/label method that uses the distance transform for their computation.

We have chosen to use parallel SIMD algorithms with boolean operations on binary images on a square lattice. The framework developed is based on several ideas: symmetries between (a) foreground and background operations, (b) 4-connected and 8-connected components, and (c) thinning and thickening operations; a minimal and most general set of 3x3 *structuring elements* (SEs) that preserve connectivity (both 4 and 8) under parallel operations; and subsets of these SEs that can operate together in parallel. For thinning, a typical goal is to find operations that preserve free ends while generating relatively smooth skeletons; for thickening, various properties such as convex hulls and exoskeleton texture are noted. The duality between thickening the foreground and thinning the background helps unify the operations; the set complement of a thickened image is an exoskeleton of the thinned background. The formalism of mathematical morphology is used because it most naturally expresses image transformations under translationally-invariant operations. The choice of 3x3 basis is pragmatic: it is the smallest allowable kernel and algorithms can be developed with smooth and conforming skeletons.

We shall see that both the choice of the SEs and their sequencing is important. Generation of a smooth skeleton, particularly with preservation of 4-connectivity, requires a delicate balance between breaking connectivity (by cascading too many different operations before updating the image) and creating a noisy, dendritic skeleton (by updating the image too frequently, leaving pixels that cannot be removed later). Examples are given that show some of the considerable variation that can result when the choice of SEs and the sequencing of the operations is altered. Rules and guidelines are presented for how operations should be sequenced to give best results.

In the derivation of the thinning and thickening algorithms, it will be necessary to distinguish between two different parallel operations. The first is the *atomic* parallel operation, where the image is thinned in parallel by matches to a specific local 3x3 pattern of ON and OFF pixels. The second is the *composite* operation, where the image is thinned in parallel by (the union of) matches to a set of local 3x3 patterns. A single *iteration* is composed of a serial sequence of four parallel operations, either *atomic* or *composite*, one for each direction.

Section 2 introduces mathematical morphology as a basis for parallel connectivity-preserving operations. The 3x3 SEs that preserve 4- and 8-connectivity on a square lattice, and are required for both thinning and thickening, are presented in Section 3. Section 4 gives results for thinning with atomic and composite parallel operations. Thickening of connected components, presented briefly in Section 5, can result in either formation of various convex hulls or growth limited only by neighboring connected regions. The paper ends with a short summary.

2 Introduction to binary morphological operations

For a survey of morphological methods, the reader is referred to the reviews of Haralick[5] or Maragos[10] for (different!) definitions of the basic operations. Our definitions are taken from Haralick[5]. Binary morphology describes translationally-invariant image-to-image operations, where the computation of each pixel in the new image is based on a set of logical operations between the pixel and some of its neighbors. The set of neighbors to be used is described by a "structuring element" (SE). The fundamental morphological operations, *erosion* and *dilation*[15], are most efficiently implemented by translating the image and either ANDing or ORing it with itself. Specifically, letting I represent the binary image and the (usually) small set S represent the *structuring element* (SE), the *erosion* $\ominus$ and *dilation* $\oplus$ of I by S are defined as

$$I \ominus S = \bigcap_{z \in S} I_{-z} \qquad (1)$$

$$I \oplus S = \bigcup_{z \in S} I_{z} \qquad (2)$$

where I_z is the *translation* of I along the pixel vector z, and the set intersection and union operations represent bitwise AND and OR, respectively. Translation is always with reference to the *center* of the SE; all 3x3 SEs used here have "centers" located at the center position. These operations can be implemented as raster operations to take advantage of the word-parallel representation of the pixels within a computer.

To handle patterns consisting of both ON and OFF pixels, Serra[15] generalized the erosion by defining a *hit-miss transform*, HMT, of an image I by a disjoint pair (A, B) of SEs as the set transformation

$$I \otimes S = I \otimes (A, B) \equiv (I \ominus A) \cap (I^c \ominus B) \tag{3}$$

where A is the *hit* SE specifying foreground pixels, B is the *miss* SE specifying background pixels, and I^c is the bit complement of I. The hit-miss SE S is in general three-valued, because it can include *don't-care* positions. The HMT returns an image with ON pixels at every location where the pattern of hits and misses matches the original image.

Simple iterative morphological operations of thinning and thickening can be described as a sequence of *atomic* parallel operations. These are defined as follows.

DEFINITION 1 *To thin an image I by a SE $S = (A, B)$, apply the HMT specified by S to I and remove any matched pixels:*

$$I \odot S \equiv I \setminus (I \otimes S) = I \setminus (I \otimes (A, B)) \tag{4}$$

where $\setminus$ denotes set subtraction

$$I \setminus J \equiv I \cap J^c \tag{5}$$

DEFINITION 2 *Likewise, to thicken the image I by S, apply the HMT and add matched pixels:*

$$I \bigcirc S \equiv I \cup (I \otimes S) = I \cup (I \otimes (A, B)) \tag{6}$$

It is easily seen that thinning and thickening of an image by a single SE are dual operations.

DEFINITION 3 *For a hit-miss SE S, denote the conjugate SE with hits and misses interchanged, by S^c:*

$$S = (A, B) \Longleftrightarrow S^c = (B, A) \tag{7}$$

Then

$$\begin{aligned} I^c \odot S = I^c \odot (A, B) = I^c \setminus (I^c \otimes (A, B)) \\ = I^c \cap (I^c \otimes (A, B))^c = (I \cup (I^c \otimes (A, B)))^c \\ = (I \cup ((I^c \ominus A) \cap (I \ominus B)))^c \\ = (I \cup (I \otimes (B, A)))^c = (I \bigcirc (B, A))^c = (I \bigcirc S^c)^c \end{aligned} \tag{8}$$

In words, *thinning the background by S is equivalent to thickening the foreground by the conjugate of S, and bit-complementing the result.*

This duality between atomic thinning and thickening operations also extends to composite thinning, using several SEs in parallel. Namely, if we take a union of HMTs with different SEs before removing or adding pixels, duality is preserved. The proof, a simple extension of the one above, is given for the case with two SEs:

$$I^c \odot (S_1, S_2) \equiv I^c \setminus ((I^c \otimes S_1) \cup (I^c \otimes S_2)) = I^c \cap ((I^c \otimes S_1)^c \cap (I^c \otimes S_2)^c)$$
$$= (I \cup ((I^c \otimes S_1) \cup (I^c \otimes S_2)))^c = (I \cup ((I \otimes S_1^c) \cup (I \otimes S_2^c)))^c \equiv (I \bigcirc (S_1^c, S_2^c))^c \tag{9}$$

These results do not depend on any special properties of the SEs used in the HMT. However, for thinning and thickening, hit-miss SEs that preserve connectivity of image components must be used. Such SEs will be defined in the next section. It should be noted that the duality of thinning and thickening operations does not imply reversability. Duality describes how the same change can be made in an image, using either thinning or thickening. But these changes are in general irreversible.

Image thinning or thickening is an iterative process, that most generally uses a set of SEs. Suppose an image is to be thinned by a set $\{S_1, S_2, \ldots, S_N\}$ of N SEs. If we simply cascade the thinning operations with respect to the set, we get the result for a single iteration:

$$I \Longrightarrow (\ldots((I \odot S_1) \odot S_2) \odot \cdots \odot S_N) \tag{10}$$

Likewise, a cascade of thickenings can be applied to an image. By duality, the thinning cascade on an image I is equivalent to the following thickening cascade on the complement of I:

$$I^c \Longrightarrow (\ldots((I^c \bigcirc S_N^c) \bigcirc S_{N-1}^c) \bigcirc \cdots \bigcirc S_1^c) \tag{11}$$

We shall see that the best thinning algorithms do not use a cascade of atomic operations. Instead, we will need to subdivide the set of N SEs into $M \leq N$ subsets $\{Z_1, Z_2, ...Z_M\}$, where each subset Z_i contains at least one SE, some of the N SEs may be contained in more than one subset, and M is typically 4, corresponding to the four lattice directions. For each subset of SEs, Z_i, the image is sequentially thinned by set subtracting the union of the HMTs specified by all SEs within the set Z_i. We refer to an operation by the union of HMTs, using the set Z_i, as a *composite* operation by Z_i. It is crucial to choose the subsets Z_i so that the composite thinning (thickening) operation does not break (join) connected components. We call such connectivity-preserving subsets Z_i of SEs *compatible*. The two major problems in devising parallel thinning or thickening algorithms can thus be stated as follows:

1. To choose an appropriate set of connectivity-preserving SEs.

2. To choose an appropriate partitioning of the set of SEs into compatible subsets Z_i for the composite operations.

The construction of SEs that conserve both 4- and 8-connectivity is discussed in the next section. The partitioning rules for thinning and thickening are given subsequently.

3 Connectivity-preserving structuring elements

A 4-connected path is described as a sequence of horizontal and vertical steps on a square lattice; an 8-connected path includes diagonal steps as well. A set of ON pixels forms an n-connected component if an n-connected path can be found between any two pixels in the set. The following definitions apply to both n-connectivity of foreground components and *dual* (12-n)-connectivity of background components. Define *weak* and *strong* connectivity preserving image operations, as follows:

DEFINITION 4 *Weak CP SE: A SE that, under an atomic operation, can alter the number of pixels in an n-connected component, but can neither split a connected component, nor join two separate components.*

We will see that this definition does not in general prevent the number of 4-connected components from changing.

DEFINITION 5 *Strong CP SE: A SE that satisfies weak CP, and additionally, under an atomic operation, neither removes all pixels within a component nor creates a new component.*

Thus, weak CP SEs are more general than strong CP SEs. A corollary of this definition is that operations using strong CP SEs preserve the number of n-connected components in the foreground and the number of dual components in the background.

These definitions explicitly emphasize symmetries between 4- and 8-connectivity SEs (i.e., between foreground and background operations). Because the more general weak CP SEs cannot create or remove 8-connected components, we could alternatively have defined weak CP SEs only for 4-connected components. However, descriptions of related phenomena (such as thinning to an endoskeleton and thickening to an exoskeleton) are much simpler using symmetric definitions.

A hit-miss SE is a set of 3-valued elements ("hit", "miss", "don't-care"). Excluding the center element, there are 3^8 3x3 hit-miss SEs. On a square lattice, any 3x3 SE is one of a set of four rotational *isomorphs*, related to each other by a sequence of 90° rotations. For thinning or thickening, these four SEs are typically used sequentially. Consider thinning from the left. We start with a hypothesis, observed to be true in practice, that all 3x3 CP SEs that can thin from the left in parallel without breaking connectivity must satisfy the template shown in Figure 1. For this template, an open circle indicates either a "miss" or a "don't care", a closed circle indicates either a "hit" or a "don't care", an empty square can be any of the three, and we ignore the center square.

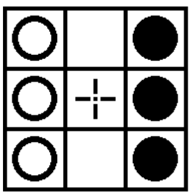

Figure 1. Most general pattern for parallel thinning from left. This is *not* a SE!

Based on this template, there are $2^6 \cdot 3^2 = 576$ possible SEs for thinning from the left. Of these, the subset that satisfies strong 4-connectivity is described by the four SEs in Figure 2. (For each SE in this paper, there are four rotational isomorphs, that describe operations from left, right, top and bottom.) For all SEs, an open circle is a "miss", a closed circle is a "hit", and an empty square is a "don't care".

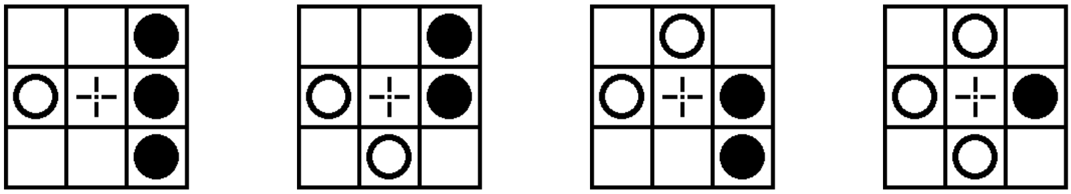

(a) S_1^4 (b) S_2^4 (c) S_3^4 (d) S_4^4

Figure 2. General SEs for strong 4-connectivity.

The analogous subset that satisfies strong 8-connectivity is shown in Figure 3.

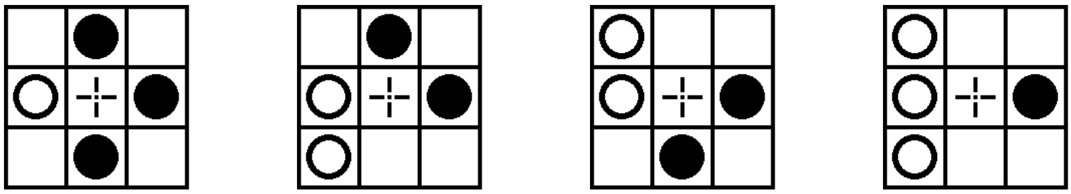

(a) S_1^8 (b) S_2^8 (c) S_3^8 (d) S_4^8

Figure 3. General SEs for strong 8-connectivity.

For weak 4 and 8-connectivity, the second and third SE of each set can be replaced by a single SE ($S_{2'}^4$ and $S_{2'}^8$). Note that $S_{2'}^4$ can remove single pixel foreground 4-connected components, and that $S_{2'}^8$ can remove single pixel background 4-connected components.

(a) $S_{2'}^4$ (b) $S_{2'}^8$

Figure 4. General SEs for weak 4- and 8-connectivity, that replace (S_2^4, S_3^4) and (S_2^8, S_3^8).

Operations that preserve 4-connectivity of foreground components also preserve 8-connectivity of background components, and v.v. This fundamental relationship between the 4- and 8-connected sets is evident from Figures 2, 3 and 4: *the SEs in each set are conjugate to each other*. Also, from these figures, it is apparent that operations that preserve 4-connectivity in the foreground will in general break 8-connectivity in the foreground, and v.v.

Consider again the duality between thinning and thickening (8). A thickening of I by one of the S^8 set is equivalent to a thinning of I^c by S^c, which is the dual of S^8 in the S^4 set; and v.v.

Figures 5 and 6 give some simple and useful specializations of the most general forms for 4-connected and 8-connected SEs, respectively. Figure 7 shows two specializations that have a high degree of symmetry and preserve *both* 4- and 8-connectivity. Use of these SEs within composite operations is illustrated in the next section.

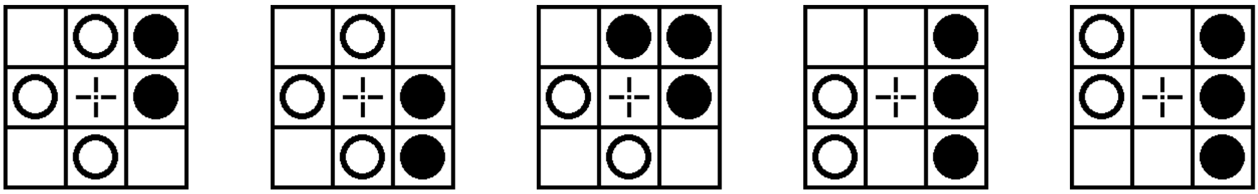

(a) S_5^4 (b) S_6^4 (c) S_7^4 (d) S_8^4 (e) S_9^4

Figure 5. Useful specialized SEs for strong 4-connectivity.

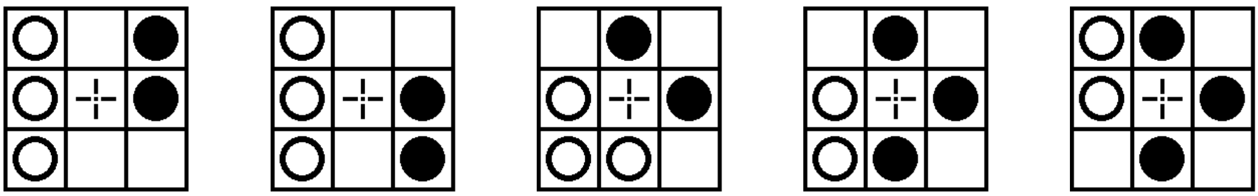

(a) S_5^8 (b) S_6^8 (c) S_7^8 (d) S_8^8 (e) S_9^8

Figure 6. Useful specialized SEs for strong 8-connectivity.

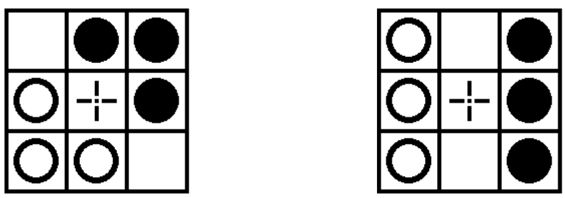

(a) $S_1^{4,8}$ (b) $S_2^{4,8}$

Figure 7. The most general SEs that preserve both 4- and 8-connectivity.

The SEs S_x^4 and S_x^8 in Figure 8 violate the basic constraints of the template in Figure 1. Although they preserve connectivity if used *sequentially* on individual pixels within the image, they break 4-connectivity (in

I and I^c, respectively) if used in parallel atomic operations.

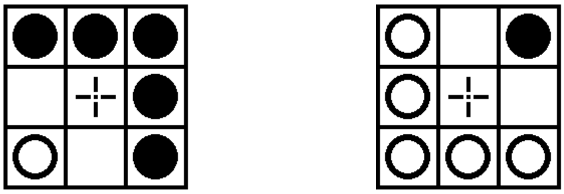

(a) S_x^4 (b) S_x^8

Figure 8. SEs that break 4-connectivity for parallel operations in I and I^c, resp.

It is useful to classify the various SEs by their symmetry properties[1]. Ordering these properties from high to low symmetry:

- Class 1. Invariant under the combination of spatial inversion and conjugation. These special SEs, $S_1^{8,4}$ and $S_2^{8,4}$, preserve both 4- and 8-connectivity.
- Class 2. The reflection about any line through the center produces a rotational isomorph (i.e., a SE that can be obtained from the first by a rotation of 90°, 180°, or 270°). For these SEs, there exists a line through the center about which the SE is invariant upon reflection. SEs with horizontal/vertical reflection symmetry (e.g., S_1^4, S_4^4, S_1^8, S_4^8) are Class 2A. Those with diagonal reflection symmetry (e.g., $S_{2'}^4$, $S_{2'}^8$, S_7^4, S_7^8) are Class 2B.
- Class 3. No reflection symmetry about any axis. There is no reflection about any line through the center that produces a rotational isomorph. These are all specializations of the most general forms that satisfy weak CP. Nevertheless, they are *very* useful for thinning.

4 Atomic and composite thinning

In this section, we consider the diversity of results of morphological thinning, examine and summarize some thinning results, and draw several general conclusions.

[1]The reader familiar with the CPT theorem of physics might note a tenuous analogy with the symmetries here. In the CPT theorem, P stands for parity (↔ spatial reflection), C for charge conjugation (↔ interchange of “hits” and “misses” in the operators), and T for time-reversal invariance (↔ addition or removal of pixels). It is believed on very general principles (and also observed) that the combination CPT is conserved in all physical processes. One might ask for the connection between the symmetries of the set of 3x3 CP operators and the analogous conservation law for connected image components!

4.1 Diversity of thinning results

Composite thinning operations, using weak and strong CP SEs, yield a variety of results depending on the specific SEs and their grouping into composite subsets. The results can be placed in six categories, ordered by generally increasing pixel removal:

1. a "blobby" result that is not completely thinned,
2. a dendritic or "noisy" skeleton,
3. a smooth skeleton, without undue erosion of free ends,
4. a smooth skeleton, with erosion of some free ends,
5. a minimal topological skeleton, or
6. a broken skeleton.

Most atomic and many composite thinning operations do not thin to completion. Define a *complete* set of SEs as one that can form a properly thinned skeleton under composite thinning applied sequentially in the four directions, as in (10). For atomic operations using the the strong and weak SEs shown in Figures 2, 3 and 4, only S_1^4 and S_2^8 (S_3^8), when used with each of their three rotational isomorphs, comprise a complete set. A noisy skeleton is formed by a complete set of SEs, but it is in a sense formed too quickly. Dendritic growth of free ends occurs spontaneously, without sufficient pruning. However, with adequate pruning of ends, a reasonably smooth skeleton can be formed without excess free end erosion. Such skeletons are desirable because they embody a simple shape representation of the connected components. Some SEs, such as S_4^4 and S_4^8 erode horizontal and vertical free ends of a thinned skeleton. This action can often be prevented by specializing to SEs such as those in Figures 5, 6, and 7. Composite operations that are able to erode both horizontal and diagonal free ends will thin to a minimal topological skeleton. Thus, a singly connected component will be reduced to a single point, a doubly connected component to a thin ring, etc. Finally, if compatible sets of SEs are not used, the connected components will be broken and may even disappear.

To preserve connectivity, it is necessary to compose the compatible sets Z_i from SEs that thin from the same "direction". The compatible sets can then be invoked sequentially either in rotation order (e.g., left, top, right, bottom) or in cross order (e.g., left, right, top, bottom). However, inspection of the SEs shows there is an ambiguity in this specification, because some SEs act to thin in a diagonal orientation. Resolution of this ambiguity (namely, the identification of compatible subsets) is a primary goal. Often, compatible subsets can be formed by combining SEs that thin from adjacent "sides". However, it is never possible to combine SEs that thin from opposite sides; this typically breaks or eliminates the skeleton.

4.2 Thinning action of SEs

Connectivity preservation for each algorithm is determined experimentally in three ways. The first step is visual inspection of the skeletons formed on a noisy scanned text image. This is usually a reliable indicator. Second, the number of 4- and 8-connected components in both foreground and background is calculated on the same image before and after thinning. Finally, the thinning algorithm is applied to an image composed of all possible 4x4 bitmaps (modulo a 90° rotation), and the number of connected components is counted before and after thinning.

Table 1 describes the action of some of the atomic thinning operations that preserve connectivity.

SE	Complete	Smoothness	Free-end erosion	Concave Hull
S_1^4	Yes	3	No	—
S_2^4	No:1	1	45°	H/V
S_4^4	No:2	N.A.	N.A.	N.A.
S_1^8	No:1	5	No	?
$S_{2'}^8$	No	4	45°	H/V
S_2^8	Yes	3	No	—
S_4^8	Yes	1	H/V	—
S_5^8	No	3	No	45°
$S_1^{4,8}$	No	2	No	H/V
$S_2^{4,8}$	No	3	No	45°

Table 1. Examples of atomic thinning operations.

In Table 1, "Concave Hull" means the orientation of unthinned segments; "No:1" means partially incomplete thinning with formation of concave hulls; "No:2" means very few pixels removed; "N.A." means not applicable because few pixels are removed; "Smoothness" of the skeleton is rated from 1 (best) to 5 (worst); "H/V" means horizontal and vertical free-end erosion or boundaries for concave hull. Ratings of skeleton smoothness are qualitatively determined from results on scanned (noisy) text images.

For 4-connected atomic thinning, only S_1^4 gives complete thinning, and is sufficient to implement a fairly dendritic approximation to a medial axis skeleton. For 8-connected atomic thinning, S_2^8 (and S_3^8) give complete thinning, but again leaving a noisy skeleton. S_4^8 has the bad combination of (1) incomplete thinning with 45° concave hulls and (2) horizontal/vertical free end erosion. Results of some of these operations are illustrated in Figure 9.

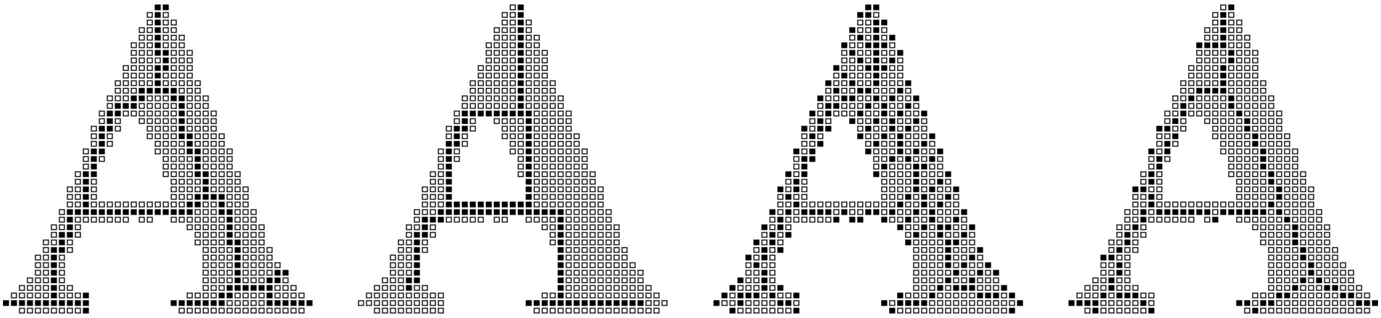

Figure 9. Atomic thinning. (a) S_1^4; (b) S_2^4; (c) S_1^8; (d) S_2^8;

Composite thinning is more interesting. Table 2 gives results for some compatible sets of SEs (i.e., sets that do not break connectivity).

SEs	Complete	Smoothness	Free-end erosion	Concave Hull
S_1^4, S_2^4, S_3^4	Yes	1	No	—
$S_1^4, S_2^4, S_3^4, S_4^4$	Yes	1	Total	—
S_1^4, S_5^4, S_6^4	Yes	1-	No	—
S_1^4, S_7^4	Yes	3	No	—
$S_1^4, S_7^4, S_{7(rot)}^4$	Yes	1	No	—
$S_1^{4,8}, S_2^{4,8}$	Yes	2	No	—
$S_1^{4,8}, S_{1(rot)}^{4,8}, S_2^{4,8}$	Yes	1	No	—
S_2^8, S_3^8	Yes	2	No	—
S_1^8, S_2^8, S_3^8	Yes	3	No	—
S_2^8, S_3^8, S_4^8	Yes	1	H/V	—
$S_2^8, S_3^8, S_2^{4,8}$	Yes	1	No	—
S_5^8, S_6^8	No:1	1	Stair	45°
S_1^8, S_5^8, S_6^8	Yes	2	No	—
$S_2^8, S_3^8, S_5^8, S_6^8$	Yes	1	No	—
$S_2^8, S_3^8, S_8^8, S_9^8$	Yes	1	No	—
$S_5^8, S_6^8, S_8^8, S_9^8$	Yes	2	No	—
$S_5^8, S_6^8, S_7^8, S_{7(rot)}^8$	Yes	1	No	—

Table 2. Examples of composite thinning operations.

In Table 2, “Total” free-end erosion means thinning to a topological minimum; “Stair” free-end erosion means removing 4-connected 45° staircases; "(rot)" indicates that a SE such as $S^{4,8}_{1(rot)}$, is rotated 90° clockwise from its partner $S^{4,8}_1$; see Table 1 for the meaning of other entries. Results of some of these operations are illustrated in Figures 10 and 11.

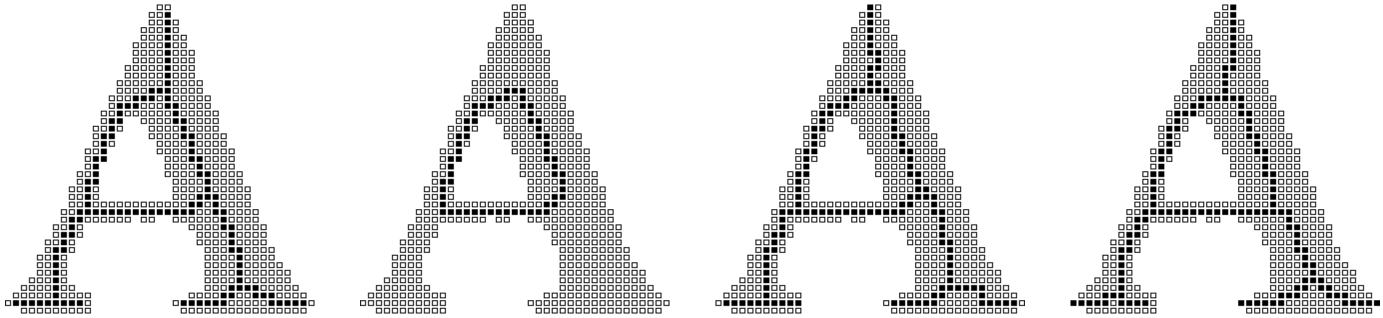

(a) (b) (c) (d)

Figure 10. 4-Connected composite thinning.
(a) S^4_1, S^4_2, S^4_3; (b) S^4_1, S^4_2, S^4_3, S^4_4; (c) S^4_1, S^4_5, S^4_6; (d) $S^{4,8}_1$, $S^{4,8}_{1(rot)}$, $S^{4,8}_2$

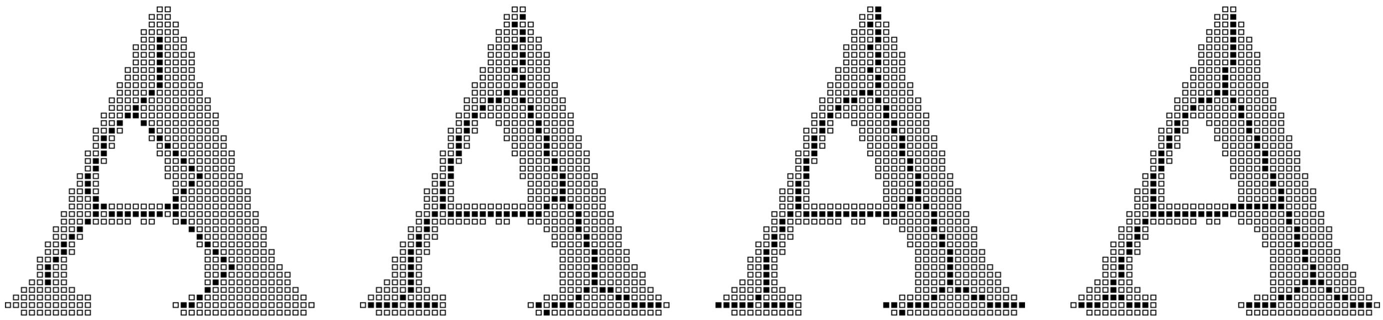

(a) (b) (c) (d)

Figure 11. 8-Connected composite thinning.
(a) S^8_5, S^8_6; (b) S^8_2, S^8_3, S^8_5, S^8_6; (c) S^8_2, S^8_3, $S^{4,8}_2$; (d) S^8_5, S^8_6, S^8_7, $S^8_{7(rot)}$

In general, the best skeletons require at least three SEs in each composite subset. As SEs are added to form the best composite sets, dendrite formation is suppressed, thinning becomes complete, and erosion of free ends is suppressed. The third factor is particularly surprising; some of the SEs *protect* the free ends from the actions of others. An example of this can be seen by comparing the results in Figures 11(a) and 11(b).

Figure 12 shows the result when a fragment of scanned text is thinned by two of the best of these algorithms. It can be seen that the 4-connected skeleton is similar in quality (smoothness, preservation of free ends) to the 8-connected one.

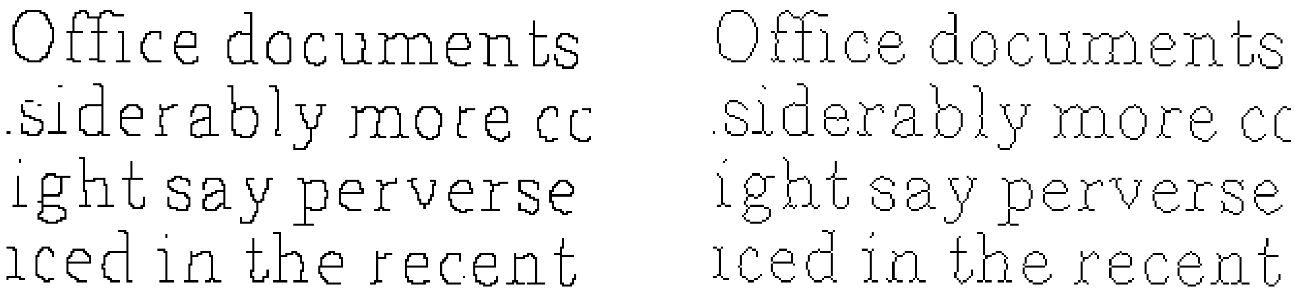

(a) (b)

Figure 12.
(a) 4-connected thining using S_1^4, S_2^4, S_3^4
(b) 8-connected thinning using S_2^8, S_3^8, S_5^8, S_6^8

The following observations can be made on compatible sets for composite thinning.

1. The very general SEs $S_{2'}^4$ and $S_{2'}^8$ that are not strong CP should *not* be used, because they tend to give poor skeletons, often broken.

2. The SEs S_4^4 and S_4^8 should not be used in combination with others because they erode horizontal and vertical free ends.

3. The order of sequential use of the four compatible sets of rotational isomorphs is not important.

4. It is advantageous to include pairs of low symmetry (Class 3) SEs such as S_2^4 and S_3^4, that are mirror reflections of each other across horizontal or vertical lines through the center. These pairs define an average thinning direction (horizontal or vertical); other SEs in the compatible sets must also thin from this average direction.

5. It is permissible to use two adjacent rotational isomorphs of those Class 1 and Class 2 SEs whose symmetry axis is on a $\pm$ 45° axis (such as $S_1^{4,8}$ and S_7^4), along with other SEs that thin horizontally or vertically from the average orientation of the rotational isomorphs. This is the only condition in which the same SE can be found in two different compatible sets Z_i.

6. The best skeletons are made using combinations of (a) low symmetry (Class 3) pairs, (b) higher symmetry (Class 1 and 2) SEs with H/V reflection symmetry, and (c) allowed rotational isomorphs.

Other combinations can be used for special purposes. For example, to thin 8-connected components to a topological minimum, one can use a combination of S_4^8 and S_x^8 to erode H/V and diagonal free-ends, respectively. (Note that S_x^8 in Figure 8b does not satisfy the general template in Figure 1; nevertheless, it preserves 8-connectivity in I).

5 Thickening

Recall that from the thinning/thickening duality, thickening I with a compatible set of SEs in S^8 is equivalent to thinning I^c with the conjugate set in S^4, and v.v. Then,

- *If a compatible set of SEs produces complete thinning to an endoskeleton, the conjugate SEs will produce complete thickening to an exoskeleton.*
- *Conversely, incomplete thinning by a compatible set of SEs is dual to thickening by the conjugate SEs to a convex hull.*

Thus, for example, we can choose SEs for 4-connected thickening to completion from compatible 8-connected sets that give complete thinning.

Self-limited convex hulls are either formed by horizontal and vertical lines, or by lines at ± 45°. However, as an algorithm for complete thickening to an exoskeleton proceeds, the freely expanding component boundaries are found to have four different shapes. These can be labelled by the slopes of the growing sides; the boundary contours between regions of constant slope do not change with expansion. Four different boundary contours have been identified: (1) 0°/90°, (2) ± 45°, (3) a right-angled quadrilateral bounded by lines with slope either ($tan^{-1}0.5$ and $-tan^{-1}2$) or ($-tan^{-1}0.5$ and $tan^{-1}2$), and (4) an octagon bounded by lines with slope $\pm tan^{-1}0.5$ and $\pm tan^{-1}2$. Table 3 gives the convex hull shapes for some self-limiting and unlimited (free expansion) thickenings; "quad" and "octagon" boundary contours refer to types (3) and (4), respectively. Table 3 does not indicate the diverse textural properties of the resulting exoskeleton.

Type	Boundary	Structuring elements
Self-limiting	H/V	S_2^8; $S_1^{4,8}$; (S_2^8, S_3^8)
Self-limiting	± 45°	S_1^4; $S_2^{4,8}$
Free expansion	H/V	S_4^4; (S_1^4, S_2^4, S_3^4); (S_4^8, $S_1^{4,8}$); (S_1^8, S_5^8, S_6^8); (S_2^8, S_3^8, S_5^8, S_6^8)
Free expansion	± 45°	S_2^4
Free expansion	quad	(S_1^4, S_7^4); ($S_1^{4,8}$, $S_2^{4,8}$)
Free expansion	octagon	(S_1^4, S_7^4, $S_{7(rot)}^4$); ($S_2^{4,8}$, $S_1^{4,8}$, $S_{1(rot)}^{4,8}$); (S_2^8, S_3^8, $S_2^{4,8}$)

Table 3. Hull and expansion shapes for some thickenings.

Thickened text images result in a variety of typographically interesting forms. Two examples with self-limiting horizontal/vertical and ± 45° convex hulls are given in Figure 13.

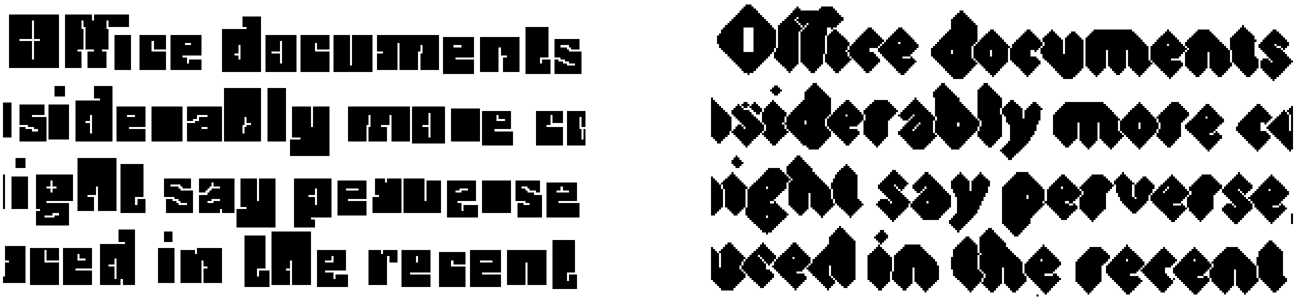

(a) (b)

Figure 13.
(a) 8-connected thickening using S_2^8 and S_3^8 (to completion)
(b) 4-connected thickening using S_1^4 (5 iterations)

6 Summary

We have explored in some depth the parallel iterative image operations that maintain component connectivity and are based on local rules with 3x3 support. The motivation is to establish rules for constructing all useful algorithms, using only logical operations, that can be carried out efficiently on either a general purpose computer or on a SIMD array processor. The 3x3 support was chosen because it is the smallest region that can be used, reasonably smooth endo- and exoskeletons can be formed, and a variety of interesting convex hulls can be produced. A number of rules, largely found experimentally, have been given in terms of the symmetry properties of strong CP SEs. Although few formal proofs are given, there is certainly a deep algebraic basis for these observations. We leave such proofs, as well as elaboration of the programme outlined in this paper, for future work. The hope is that questions have been posed in such a way as to inspire and perhaps even direct further inquiry.

We have constructed the least restrictive 3x3 hit-miss SEs that can be used morphologically to preserve either 4-connected or 8-connected regions of binary images. From these SEs a few less general but very useful pairs of SEs have been derived. The SEs vary in the degree to which they erode and smooth the skeleton. Nevertheless, many combinations of these SEs have been found that leave reasonably smooth approximations to a medial axis skeleton, for both 4-connected and 8-connected skeletons, without undue erosion of skeletal end points. This is particularly encouraging for 4-connected skeletons, for which prevention of dendritic growth has been problematic. High symmetry SEs can be used in parallel to preserve both 4- and 8-connectivity.

Because of the duality between thinning and thickening, results with parallel composite thinning can be immediately extended to thickening with conjugate SEs. With thickening we naturally focus on properties such as convex hulls and aesthetics of partial and completed operations. Notwithstanding the low degree of symmetry of the square lattice, there are several parallel unbounded thickening operations with an 8-sided expanding hull. *Regularized* images, which can be formed by sequentially thinning to a skeleton and thickening by a fixed amount, may be useful for some aspects of image analysis.

References

[1] C. Arcelli, L. Cordella and S. Levialdi, "Parallel thinning of binary pictures", *Electronic Letters* **11**, 148-149 (1975)

[2] H. Blum, "A transformation for extracting new descriptors of shape", in *Models for the perception of Speech and Visual Forms*, W. Wathen-Dunn, ed, M.I.T. Press, Cambridge, MA, 1967.

[3] R. T. Chin, H-K Wan, D. L. Stover and R. D. Iverson, "A one-pass thinning algorithm and its parallel implementation", *CVGIP* **40**, 30-40 (1987)

[4] M. J. E. Golay, "Hexagonal parallel pattern transformations", *IEEE Trans. Comput.* **C-18**, 733-740 (1969).

[5] R. M. Haralick, S. R. Sternberg and X. Zhuang, "Image analysis using mathematical morphology", *IEEE Trans. on Pattern Analysis and Machine Intelligence* **PAMI-9**, 532-550 (1987).

[6] R. W. Hall, "Fast parallel thinning algorithms: parallel speed and connectivity preservation", *Comm. ACM* **32**, 124-131 (1989).

[7] C. M. Holt, A. Stewart, M. Clint, and R. H. Perrott, "An improved parallel thinning algorithm", *Comm. ACM* **30**, 156-160 (1987).

[8] H. E. Lu and P. S. P. Wang, "A comment on 'A fast parallel algorithm for thinning digital patterns", *Comm. ACM* **29**, 239-242 (1986).

[9] P. A. Maragos and R. W. Schafer, "Morphological skeleton representation and coding of binary images", *IEEE Trans. on Acoustics, Speech, and Signal Processing* **ASSP-34**, 1228-1244 (1986).

[10] P. A. Maragos, "Tutorial on advances in morphological image processing and analysis", *Optical Engineering* **26**, 623-632 (1987).

[11] A. Rosenfeld, "Connectivity in digital pictures", *J. ACM* **17**, 146-160 (1970).

[12] A. Rosenfeld, "A characterization of parallel thinning algorithms", *Information and Control* **29**, 286-291 (1975).

[13] A. Rosenfeld and L. S. Davis, "A note on thinning", *IEEE Trans. Sys., Man, Cybern.* **6**, 226-228 (1976).

[14] A. Rosenfeld and A. C. Kak, *Digital Picture Processing, Vol. 2*, Academic Press, London, 1982. pp. 232-240.

[15] J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, London, 1982.

[16] R. Stefanelli and A. Rosenfeld, “Some parallel thinning algorithms for digital pictures”, *J. ACM* **18**, 255-264 (1971).

[17] L. Vincent, “Efficient computation of various types of skeletons”, *Proc. SPIE Conf. Medical Imaging V*, San Jose, CA, Februrary 1991.

[18] P. S. P. Wang and Y. Y. Zhang, “A fast serial and parallel thinning algorithm”, in *Proceedings of the Eighth European Meeting on Cybernetics and Systems Research 86*, Vienna, Austria, 909-915 (1986).

[19] T. Y. Zhang and C. Y. Suen, “A fast parallel algorithm for thinning digital patterns”, *Comm. ACM* **27**, 236-239 (1984).