Geometry of digital spaces

G. T. Herman
Medical Image Processing Group, Department of Radiology
University of Pennsylvania, Blockley Hall, Fourth Floor
423 Guardian Drive, Philadelphia, PA 19104–6021, U.S.A.

ABSTRACT

In the recently published Ref. 1 the author surveys a number of aspects of the geometry of digital spaces. In this article we exemplify the approach of that book, by providing a self-contained proof of one of its final result which is to do with the correctness and the characterization of the output of a general purpose boundary-tracking algorithm.

Keywords: Digital geometry, Digital topology, Digital spaces, Boundary tracking, Surfaces, Hyper-cubes

1. SURFACES IN DIGITAL SPACES

Let M be any set, and ρ be a *binary relation* on M (i.e., ρ is a subset of M^2, the set of all ordered pairs of elements of M). If $(c,d) \in \rho$, then we say that c is *ρ-adjacent to d* and that d is *ρ-adjacent from c* and, in case ρ is a *symmetric relation* (meaning that $(c,d) \in \rho$ if, and only if, $(d,c) \in \rho$), that c and d are *ρ-adjacent*. We will often use the word "adjacency" to refer to a symmetric binary relation.

We use Z to denote the set of all integers. For any positive integer N, we define

$$Z^N = \{\, (c_1, \cdots, c_N) \,|\, c_n \in Z, \ \text{ for } 1 \leq n \leq N \} \,. \tag{1.1}$$

In the case when $M = Z^N$, we will be repeatedly dealing with two binary relations, ω_N and δ_N, defined as:

$$(c,d) \in \omega_N \Leftrightarrow \sum_{n=1}^{N} |c_n - d_n| = 1 \,. \tag{1.2}$$

$$(c,d) \in \delta_N \Leftrightarrow 1 \leq \sum_{n=1}^{N} (c_n - d_n)^2 \leq 2 \,. \tag{1.3}$$

Note that, for any positive integer N, $\omega_N \subset \delta_N$. ($A \subset B$ denotes the phrase "A is a subset of B.")

Let A be a subset of M. For any c and d in A, the sequence $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ of elements of A is said to be a *ρ-path in A connecting c to d*, if $c^{(0)} = c$, $c^{(K)} = d$ and, for $1 \leq k \leq K$, $c^{(k-1)}$ is ρ-adjacent to $c^{(k)}$. We call K the *length* of this path. Note that the length is measured not by the number of elements in the path, but rather by the number of steps needed to get from its beginning to its end. In particular, there are ρ-paths of length zero (such as $\langle c \rangle$); we refer to them as *trivial paths*. If there is a ρ-path in A connecting c to d, then we say that c is *ρ-connected in A to d*. (Since every element of A is ρ-connected in A to itself by a trivial path, ρ-connectedness in A is a reflexive relation) A nonempty subset A of M is said to be a *ρ-connected subset* if, for any c and d in A, c is ρ-connected in A to d.

If there are ρ-paths in A connecting c to d and d to e, then they can be combined into a ρ-path in A connecting c to e, hence ρ-connectedness in A is also a transitive relation If it is also the case that ρ-connectedness in A is a symmetric relation on A (and hence an *equivalence relation* on A), then it partitions A into *ρ-components* (nonempty ρ-connected subsets which are not proper subsets of any other ρ-connected subset of A). If ρ happens to be symmetric, then ρ-connectedness in A is guaranteed to be a symmetric (and hence an equivalence) relation on A.

E-mail: gabor@mipg.upenn.edu

Part of the SPIE Conference on Vision Geometry VII • San Diego, California • July 1998
SPIE Vol. 3454 • 0277-786X/98/$10.00

In the special case when $A = M$, we use the phrases *ρ-path* and *ρ-connected* instead of ρ-path in A and ρ-connected in A. For any binary relation ρ on a set M, we define its *transitive closure* as the binary relation ρ^* on M, defined by: $(c,d) \in \rho^*$ if, and only if, c is ρ-connected to d. As we have already pointed out, ρ^* is always a reflexive and transitive relation on M. If, in addition, ρ^* also happens to be symmetric, then it is an equivalence relation and so partitions M into ρ-components.

A *digital space* is a pair (V, π), where V is an arbitrary nonempty set and π is a symmetric binary relation on V such that V is π-connected. Sometimes we will refer to π as the *proto-adjacency* of the digital space (V, π). Elements of V will be called *spels* (short for spatial elements) and elements of π will be called *surfels* (short for surface elements). Any nonempty subset of π will be called a *surface* in (V, π). Occasionally we will need to restrict our attention to *antisymmetric surfaces* S, meaning that if $(c,d) \in S$, then $(d,c) \notin S$.

We will be particularly concerned with a special class of surfaces, which we refer to as boundaries. Let (V, π) be a digital space and let O and Q be subsets of V. Then the *boundary in* (V, π) *between* O and Q is defined to be

$$\partial(O,Q) = \{ (c,d) \mid (c,d) \in \pi,\ c \in O \ \&\ d \in Q \} . \tag{1.4}$$

If it is not empty, then a boundary in (V, π) is a surface in (V, π). In what follows, we will normally be dealing with nonempty boundaries between disjoint sets. Clearly, such boundaries are antisymmetric surfaces.

Let (V, π) be a digital space and let S be a surface in it. We define the *immediate interior* $II(S)$ and the *immediate exterior* $IE(S)$ of S as follows:

$$\begin{aligned} II(S) &= \{ c \mid (c,d) \in S \text{ for some } d \text{ in } V \} , \\ IE(S) &= \{ d \mid (c,d) \in S \text{ for some } c \text{ in } V \} . \end{aligned} \tag{1.5}$$

We say that a π-path $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ *crosses* the surface S if there is a k, $1 \le k \le K$, such that either $(c^{(k-1)}, c^{(k)}) \in S$ or $(c^{(k)}, c^{(k-1)}) \in S$. The *interior* $I(S)$ and the *exterior* $E(S)$ of S are defined as follows:

$$\begin{aligned} I(S) &= \{ c \in V \mid \text{there exists a } \pi\text{-path connecting } c \text{ to an element of } II(S) \text{ which does not cross } S \} , \\ E(S) &= \{ c \in V \mid \text{there exists a } \pi\text{-path connecting } c \text{ to an element of } IE(S) \text{ which does not cross } S \} . \end{aligned} \tag{1.6}$$

A surface S in a digital space (V, π) is said to be *near-Jordan* if every π-path from any element of $II(S)$ to any element of $IE(S)$ crosses S. (This immediately implies that a near-Jordan surface is an antisymmetric surface.) We remark that if S is not a near-Jordan surface, then the intersection of its interior and its exterior is necessarily nonempty. (This is because there is a path, not crossing S, from its immediate interior to its immediate exterior; all the spels in this path are in both the interior and the exterior of S, as can be seen from (1.6), the symmetry of proto-adjacency and the symmetrical definition of "crosses.")

Lemma 1.1. Let S be a surface in a digital space (V, π). Then the following three conditions are equivalent.

(i) S is near-Jordan.

(ii) Every π-path from any element of $I(S)$ to any element of $E(S)$ crosses S.

(iii) $I(S) \cap E(S) = \emptyset$. (As usual, $\emptyset$ denotes the empty set.)

Furthermore, if these conditions are satisfied, then it is also the case that

$$S = \partial(I(S), E(S)) . \tag{1.7}$$

Proof. Let S be any surface in a digital space (V, π). Suppose there is a π-path $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ not crossing S connecting a c in $I(S)$ to a d in $E(S)$. By (1.6), there is also a π-path $\langle e^{(0)}, \cdots, e^{(L)} \rangle$ not crossing S connecting c to an element of $II(S)$ and a π-path $\langle d^{(0)}, \cdots, d^{(M)} \rangle$ not crossing S connecting d to an element of $IE(S)$. Then $\langle e^{(L)}, \cdots, e^{(0)} = c = c^{(0)}, \cdots, c^{(K)} = d = d^{(0)}, \cdots, d^{(M)} \rangle$ is a π-path not crossing S which connects an element of $II(S)$ to an element of $IE(S)$. Therefore, by definition, S is not near-Jordan. This argument shows that (i) implies (ii). If $I(S)$ and $E(S)$ have an element c in common, then the trivial π-path $\langle c \rangle$ is from an element of $I(S)$ to an element of $E(S)$ and does not cross S. This shows that (ii) implies (iii). That (iii) implies (i) was essentially proven in the paragraph preceding the statement of the lemma. Finally, it is the case for any surface S that it is a subset of $\partial(I(S), E(S))$. (This follows trivially from the definitions of immediate interior and exterior and of interior and exterior.) Suppose now that (i) - (iii) hold and that the surfel (c, d) is in $\partial(I(S), E(S))$. Then d belongs to $E(S)$, and so by (iii) d does not belong to $I(S)$; hence (d, c) is not in S. The π-path $\langle c, d \rangle$ is from an element of $I(S)$ to an element of $E(S)$ and so, as condition (ii) is satisfied, it crosses S. Hence (c, d) is in S. This proves that any of the conditions (i) - (iii) implies (1.7). □

2. SIMPLY CONNECTEDNESS

Let S be a surface in a digital space (V,π). For any practical application, it would be impossible to determine whether S is near-Jordan by examining all π-paths from $II(S)$ to $IE(S)$. It is desirable to have a result which says that S is near-Jordan if some local condition is satisfied at every surfel of S.

In classical topology, a simply connected space is (intuitively speaking) a connected space in which every loop can be continuously pulled to a point without leaving the space. There is an infinity of corresponding notions for digital spaces: for every positive integer D (reflecting how large a digital step is allowed to replace the notion of continuity), there is a class of D-simply connected digital spaces, whose definition now follows.

If

$$P = \langle c^{(1)}, \cdots, c^{(m)}, d^{(0)}, \cdots, d^{(n)}, e^{(1)}, \cdots, e^{(l)} \rangle \tag{2.1}$$

and

$$P' = \langle c^{(1)}, \cdots, c^{(m)}, f^{(0)}, \cdots, f^{(k)}, e^{(1)}, \cdots, e^{(l)} \rangle \tag{2.2}$$

are π-paths in the digital space (V,π), such that

$$f^{(0)} = d^{(0)}, \quad f^{(k)} = d^{(n)}, \quad \text{and} \quad 1 \le k+n \le D+2, \tag{2.3}$$

then P and P' are said to be *elementarily D-equivalent.* (In this definition, m or l or both in (2.1) may be zero.) Two π-paths, P and P' in a digital space (V,π) are said to be *D-equivalent,* if there is a sequence of π-paths $P_0, \cdots, P_L$ $(L \ge 0)$ in the digital space, such that $P_0 = P$, $P_L = P'$ and, for $1 \le l \le L$, P_{l-1} and P_l are elementarily D-equivalent.

We demonstrate 2-equivalent π-paths in (Z^2,ω_2) in Figure 2.1: $\langle a,b,c,d,e,f,g,h,i,j,k,l,a\rangle$ is elementarily 2-equivalent to $\langle a,b,c,d,e,d,g,h,i,j,k,l,a\rangle$ by substituting in (2.1) $n=2$, $d^{(0)}=e$, $d^{(1)}=f$, $d^{(2)}=g$ and in (2.2) $k=2$, $f^{(0)}=e$, $f^{(1)}=d$, $f^{(2)}=g$; $\langle a,b,c,d,e,d,g,h,i,j,k,l,a\rangle$ is elementarily 2-equivalent to $\langle a,b,c,d,g,h,i,j,k,l,a\rangle$ by substituting in (2.1) $n=2$, $d^{(0)}=d$, $d^{(1)}=e$, $d^{(2)}=d$ and in (2.2) $k=0$, $f^{(0)}=d$.

A *loop* (*of length* K) in a digital space (V,π) is a π-path $\langle c^{(0)}, \cdots, c^{(K)}\rangle$ such that $c^{(K)} = c^{(0)}$. In particular, for any spel c, $\langle c\rangle$ is a loop, and is called a *trivial loop.* (The notion of a trivial loop and the previously defined notion of a trivial path are identical; we select the terminology to be used depending on the context.) We note that any loop of length 1, 2, or 3 is automatically D-equivalent to a trivial loop, for any positive integer D. A digital space is said to be *D-simply connected* if every loop in the digital space is D-equivalent to a trivial loop.

Let S be a surface in a digital space (V,π). We say that a surfel (c,d) in (V,π) *crosses* S if exactly one of $(c,d)\in S$ or $(d,c)\in S$ is true. Let $P = \langle c^{(0)}, \cdots, c^{(K)}\rangle$ be a π-path in (V,π). We say that the *crossing parity* $p_S P$ *of* P *through* S is *even* (or zero; i.e., $p_S P = 0$) if the number of surfels among $(c^{(0)}, c^{(1)}), \cdots, (c^{(K-1)}, c^{(K)})$ that cross S is even and we say that it is *odd* (or one; i.e., $p_S P = 1$) if this number is odd. We use the notation $\oplus$ for *modulo 2 addition* of parities (i.e.; $0\oplus 0 = 1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$).

It is easy to see that, for any surface S in a digital space, the crossing parity through S is even for any loop in the digital space whose length is not greater than two. Also, cyclic permutation of a loop does not influence its crossing parity through a surface S, since, for $1 \le k \le K$,

$$p_S\langle c^{(0)}, \cdots, c^{(k-1)}, c^{(k)}, \cdots, c^{(K)}\rangle = p_S\langle c^{(k)}, \cdots, c^{(K)} = c^{(0)}, \cdots, c^{(k-1)}, c^{(k)}\rangle\,. \tag{2.4}$$

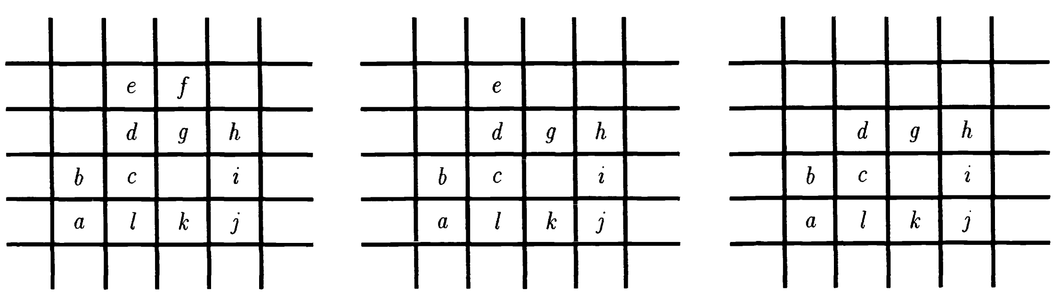

Figure 2.1. Demonstration of equivalence of π-paths in (Z^2,ω_2): $\langle a,b,c,d,e,f,g,h,i,j,k,l,a\rangle$ is 2-equivalent to $\langle a,b,c,d,g,h,i,j,k,l,a\rangle$, since it is elementarily 2-equivalent to $\langle a,b,c,d,e,d,g,h,i,j,k,l,a\rangle$, which is elementarily 2-equivalent to $\langle a,b,c,d,g,h,i,j,k,l,a\rangle$.

In addition, reversing a π-path does not change its crossing parity. It is also easy to see that if $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ and $\langle d^{(0)}, \cdots, d^{(L)} \rangle$ are π-paths such that $c^{(K)} = d^{(0)}$, then

$$p_S\langle c^{(0)}, \cdots, c^{(K)}, d^{(1)}, \cdots, d^{(L)} \rangle = p_S\langle c^{(0)}, \cdots, c^{(K)} \rangle \oplus p_S\langle d^{(0)}, \cdots, d^{(L)} \rangle . \tag{2.5}$$

Theorem 2.1. If S is a near-Jordan surface in a digital space, then the crossing parity through S is odd for any π-path $P = \langle c^{(0)}, \cdots, c^{(K)} \rangle$ such that $(c^{(0)}, c^{(K)}) \in S$.

Proof. First note that (according to Lemma 1.1), since S is near-Jordan, $I(S) \cap E(S) = \emptyset$ and $S = \partial(I(S), E(S))$. We prove by induction that, for any $0 \leq k \leq K$,

$$p_S\langle c^{(0)}, \cdots, c^{(k)} \rangle = \begin{cases} 0, & \text{if } c^{(k)} \in I(S), \\ 1, & \text{otherwise.} \end{cases} \tag{2.6}$$

Since $c^{(K)} \in E(S)$, this is sufficient to prove the theorem.

Clearly, (2.6) is true for $k = 0$. Suppose that (2.6) is true for some $k - 1$, where $1 \leq k \leq K$. We prove that it is also true for k. We make repeated use of the following special case of (2.5):

$$p_S\langle c^{(0)}, \cdots, c^{(k-1)}, c^{(k)} \rangle = p_S\langle c^{(0)}, \cdots, c^{(k-1)} \rangle \oplus p_S\langle c^{(k-1)}, c^{(k)} \rangle . \tag{2.7}$$

In case $c^{(k-1)} \in I(S)$, the first term on the right hand side of (2.7) is 0 and the second term is 0 if $c^{(k)} \in I(S)$ and 1 otherwise. In case $c^{(k-1)} \notin I(S)$, the first term on the right hand side of (2.7) is 1 and the second term is 1 if $c^{(k)} \in I(S)$ and 0 otherwise. In either case, (2.6) is true for k. □

A surface S in a digital space (V, π) is said to be *D-locally-Jordan* (where D is a positive integer) if $p_S\langle c^{(0)}, \cdots, c^{(K)} \rangle$ is odd for any π-path such that $(c^{(0)}, c^{(K)}) \in S$ and $2 \leq K \leq D + 1$. By Theorem 2.1, if a surface S in a digital space (V, π) is near-Jordan, then it is D-locally-Jordan for all positive D.

Lemma 2.2. Any loop of length not more than $D+2$ has even crossing parity through any D-locally-Jordan surface in any digital space (V, π).

Proof. We have already pointed out that the crossing parity through any surface is even for any loop of length not more than two. Let S be an D-locally-Jordan surface. Consider a loop $L = \langle c^{(0)}, \cdots, c^{(K)} \rangle$ with $3 \leq K \leq D + 2$. If there does not exist a k, $1 \leq k \leq K$, such that $(c^{(k-1)}, c^{(k)})$ crosses S, then we are done. Otherwise, for such a k, exactly one of $(c^{(k-1)}, c^{(k)}) \in S$ or $(c^{(k)}, c^{(k-1)}) \in S$ is true. If $(c^{(k)}, c^{(k-1)}) \in S$, $P = \langle c^{(k)}, \cdots, c^{(K)} = c^{(0)}, \cdots, c^{(k-1)} \rangle$ is a π-path of length $K - 1$ with $2 \leq K - 1 \leq D + 1$ and so the fact that S is D-locally-Jordan implies $p_S P = 1$. By application of (2.4) and (2.5), we have

$$p_S L = p_S P \oplus p_S\langle c^{(k-1)}, c^{(k)} \rangle = 1 \oplus 1 = 0 . \tag{2.8}$$

If $(c^{(k-1)}, c^{(k)}) \in S$, a similar argument, which also makes use of the fact that reversing a π-path does not change its crossing parity, can be used to derive the same conclusion. □

Lemma 2.3. Let S be an D-locally-Jordan surface in a digital space (V, π). If P and P' are D-equivalent π-paths, then they have the same crossing parity through S.

Proof. By the definition of D-equivalent, it is sufficient to prove that if P and P' satisfy (2.1), (2.2), and (2.3), then they have the same crossing parity through S.

By applying (2.5), we get

$$p_S P = p_S\langle c^{(1)}, \cdots, c^{(m)}, d^{(0)} \rangle \oplus p_S\langle d^{(0)}, \cdots, d^{(n)} \rangle \oplus p_S\langle d^{(n)}, e^{(1)}, \cdots, e^{(l)} \rangle \tag{2.9}$$

and

$$p_S P' = p_S\langle c^{(1)}, \cdots, c^{(m)}, f^{(0)} \rangle \oplus p_S\langle f^{(0)}, \cdots, f^{(k)} \rangle \oplus p_S\langle f^{(k)}, e^{(1)}, \cdots, e^{(l)} \rangle . \tag{2.10}$$

Therefore, using (2.3), the invariance of crossing parity under reversal, and (2.5), we get

$$\begin{aligned} p_S P \oplus p_S P' &= p_S\langle d^{(0)}, \cdots, d^{(n)} \rangle \oplus p_S\langle f^{(0)}, \cdots, f^{(k)} \rangle \\ &= p_S\langle d^{(0)}, \cdots, d^{(n)} = f^{(k)}, \cdots, f^{(0)} = d^{(0)} \rangle = 0 . \end{aligned} \tag{2.11}$$

The last equality follows from the previous lemma combined with (2.3). □

The results up to now apply to digital spaces which do not have to be D-simply connected for any D. The next lemma makes essential use of D-simple connectedness.

Lemma 2.4. If S is an D-locally-Jordan antisymmetric surface in an D-simply connected digital space (V, π), then S is near-Jordan if either (and hence both) of the following two equivalent conditions holds.

(i) For any $c \in II(S)$ and $d \in II(S)$, there exists a π-path P from c to d such that $p_S P = 0$.

(ii) For any $c \in IE(S)$ and $d \in IE(S)$, there exists a π-path P from c to d such that $p_S P = 0$.

Proof. Evidently the two conditions are equivalent. Indeed, if $c \in II(S)$ and $d \in II(S)$, then there exist $c' \in IE(S)$ and $d' \in IE(S)$ such that $(c, c') \in S$ and $(d, d') \in S$; hence if there exists a π-path of even crossing parity from c' to d', then there also exists one from c to d, so that (ii) implies (i). Similarly, (i) implies (ii).

In what follows, we prove that S is near-Jordan if (ii) holds. We do this by supposing that S is not near-Jordan and showing that this, together with (ii), leads to a contradiction.

First, we show that there exists a π-path $P_1 = \langle c^{(1)}, \cdots, c^{(K)} \rangle$ such that $c^{(1)} \in II(S)$, $c^{(K)} \in IE(S)$ and $p_S P_1 = 0$. Indeed, since S is supposed to be not near-Jordan, there is a π-path from $II(S)$ to $IE(S)$ that does not cross S. Clearly, any such π-path has the required properties.

Next, we show that there exists a π-path $P_3 = \langle e^{(1)}, \cdots, e^{(L)} \rangle$ such that $(e^{(1)}, e^{(L)}) \in S$ and $p_S P_3 = 0$. Let P_1 be the π-path of the last paragraph. Let $c^{(0)}$ be such that $(c^{(1)}, c^{(0)}) \in S$. By (ii), there exists a π-path $P_2 = \langle c^{(K)} = d^{(0)}, \cdots, d^{(L)} = c^{(0)} \rangle$ from $c^{(K)}$ to $c^{(0)}$ such that $p_S P_2 = 0$. Then $P_3 = \langle c^{(1)}, \cdots, c^{(K)}, d^{(1)}, \cdots, d^{(L)} \rangle$ is a π-path from $c^{(1)}$ to $d^{(L)}$ such that $(c^{(1)}, d^{(L)}) \in S$ and, by (2.5), $p_S P_3 = p_S P_1 \oplus p_S P_2 = 0$.

For a π-path P_3 satisfying the properties listed at the beginning of the previous paragraph, let $e^{(0)} = e^{(L)}$. By the antisymmetry of S, $p_S \langle e^{(0)}, e^{(1)} \rangle = 1$ and so, by (2.5), $P_4 = \langle e^{(0)}, e^{(1)}, \cdots, e^{(L)} \rangle$ is a loop such that $p_S P_4 = p_S \langle e^{(0)}, e^{(1)} \rangle \oplus p_S P_3 = 1$. Since (V, π) is D-simply connected, P_4 is D-equivalent to a trivial loop, whose crossing parity through S is zero. Since S is D-locally-Jordan, according to Lemma 2.3, we must also have $p_S P_4 = 0$, contradicting the fact that $p_S P_4 = 1$. □

Theorem 2.5. For any positive integer N, (Z^N, ω_N) is 2-simply connected.

Proof. We show that every loop in (Z^N, ω_N) is 2-equivalent to a trivial loop. We do this by induction on the length of the loop. We have already noted that, in general, every loop of length 1 or 2 is 2-equivalent to a trivial loop. Suppose that every loop in (Z^N, ω_N) of length less than some $K > 2$ is 2-equivalent to a trivial loop. Consider a loop $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ in (Z^N, ω_N) of length K. We now show that it is 2-equivalent to a loop of length $K - 1$ or $K - 2$ and thus, by the induction hypothesis, is 2-equivalent to a trivial loop.

Since $c^{(1)}$ is ω_N-adjacent to $c^{(0)}$, we have $c^{(1)} \neq c^{(0)}$. (To illustrate our argument, consider Figure 2.1. In that figure $c^{(0)} = a = (0,0)$ and $c^{(1)} = b = (0,1)$.) Then there is a unique j such that $c_j^{(1)} \neq c_j^{(0)}$. Without loss of generality, assume that $c_j^{(1)} > c_j^{(0)}$. (In Figure 2.1, $j = 2$.) Let $z = \max_{1 \leq k \leq K} \left\{ c_j^{(k)} \right\}$. (In Figure 2.1, $z = 3$.) Let l be the largest integer in the range $0 < l < K$ such that $c_j^{(l)} = z$. (In Figure 2.1, in the left column $l = 5$ with $c^{(5)} = f = (2,3)$ and in the middle column $l = 4$ with $c^{(4)} = e = (1,3)$.) Clearly, $c_j^{(l+1)} = z - 1$. (In Figure 2.1, in the left column $c^{(l+1)} = c^{(6)} = g = (2,2)$ and in the middle column $c^{(l+1)} = c^{(5)} = d = (1,2)$.) Let k be the smallest integer such that, for all i in the range $0 < k \leq i \leq l < K$, we have $c_j^{(i)} = z$. (In Figure 2.1, $k = 4$ in both the left and the middle column.) Clearly, $c_j^{(k-1)} = z - 1$. (In Figure 2.1, $c^{(k-1)} = c^{(3)} = d = (1,2)$ in both the left and middle columns.) We will now use induction on $l - k$.

If $l - k = 0$, then $k = l$ and therefore $c^{(k-1)} = c^{(k+1)}$. (This case is illustrated in Figure 2.1 by the middle column, for which $k = l = 4$ and $c^{(k-1)} = c^{(k+1)} = d$.) In this case the loop

$$\langle c^{(0)}, \cdots, c^{(k-2)}, c^{(k-1)}, c^{(k)}, c^{(k+1)}, c^{(k+2)}, \cdots, c^{(K)} \rangle \tag{2.12}$$

is elementarily 2-equivalent to the loop

$$\langle c^{(0)}, \cdots, c^{(k-2)}, c^{(k-1)} = c^{(k+1)}, c^{(k+2)}, \cdots, c^{(K)} \rangle \tag{2.13}$$

which has length $K - 2$ and we are done. (The loop in (2.12) is illustrated by the loop of the middle column of Figure 2.1, while the loop in (2.13) is illustrated by the loop of the right column of Figure 2.1.)

Suppose now (induction hypothesis) that whenever $l - k = h$, then $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ is 2-equivalent to a loop of length $K - 1$ or $K - 2$. We now show that the same conclusion holds if $l - k = h + 1$. Let

$$\langle c^{(0)}, c^{(1)}, \cdots, c^{(k)}, \cdots, c^{(l-1)}, c^{(l)}, c^{(l+1)}, \cdots, c^{(K)} \rangle \tag{2.14}$$

be a loop. Define j, z, l, and k for this loop as above and suppose $l - k = h + 1$. (This is the case in Figure 2.1 for the loop associated with the left column with $l = 5$, $k = 4$ and, hence, $h = 0$. Note also that for this loop $j = 2$ and $z = 3$.) Let $c'^{(l)}$ be a spel such that, for $1 \le n \le N$,

$$c'^{(l)}_n = \begin{cases} z - 1, & \text{if } n = j, \\ c^{(l-1)}_n, & \text{otherwise.} \end{cases} \tag{2.15}$$

(In the left column of Figure 2.1, $c'^{(l)} = c'^{(5)} = (1,2) = d$.) Clearly, $c'^{(l)}$ is proto-adjacent to $c^{(l-1)}$. Also, $c^{(l+1)}$ differs from $c^{(l)}$ in exactly the jth component (which is $z - 1$ for the former) and $c^{(l)}$ differs from $c^{(l-1)}$ in exactly one component which is other than the jth; thus $c'^{(l)}$ is proto-adjacent to $c^{(l+1)}$. (In the left column of Figure 2.1, $c^{(l-1)} = c^{(4)} = e$ and $c^{(l+1)} = c^{(6)} = g$, both of which are proto-adjacent to $c'^{(l)} = c'^{(5)} = d$.) It follows that

$$\langle c^{(0)}, c^{(1)}, \cdots, c^{(k)}, \cdots, c^{(l-1)}, c'^{(l)}, c^{(l+1)}, \cdots, c^{(K)} \rangle \tag{2.16}$$

is also a loop and is easily seen to be elementarily 2-equivalent to the loop in (2.14). (The loop in (2.14) is illustrated by the loop of the left column of Figure 2.1, while the loop in (2.16) is illustrated by the loop of the middle column of Figure 2.1.) Furthermore, it follows from (2.15) that for the loop in (2.16) the condition of the induction hypothesis holds. (Indeed, we have already seen that for the middle column of Figure 2.1 we have $l - k = 0 = h$.) So, by the induction hypothesis, the loop in (2.16) is 2-equivalent to a loop of length $K - 1$ or $K - 2$. From this it follows that the loop in (2.14) is also 2-equivalent to a loop of length $K - 1$ or $K - 2$. □

3. GENERAL BOUNDARY TRACKING

We define a *digital picture* over the digital space (V, π) to be a triple (V, π, f), where f is a function whose domain is V. The choice of topics in this book is biased towards the study of *binary pictures*, which are defined to be digital pictures in which the range of f is the two-element set $\{0, 1\}$. We refer to those spels which map into 0 under f as *0-spels* and we refer to those spels which map into 1 under f as *1-spels*.

If O is the set of all 1-spels and Q is the set of all 0-spels in a binary picture (V, π, f), then we refer to elements of $B = \partial(O, Q)$ as *bels* (short for boundary elements) in (V, π, f). We say that the binary relation λ on B is a *bel-adjacency* if

(i) λ^* is symmetric and
(ii) for any b in B, there is a finite number (denoted by $in_\lambda(b)$ and called the *indegree* of the bel b) of different elements a in B for which $(a, b) \in \lambda$.

We propose an algorithm for the following task: **Given** a binary picture (V, π, f), a bel-adjacency λ, and a bel o, such that the λ-component S of B which contains o is finite, **find** S.

General Bel-Tracking Algorithm

(1) Put o into L and S and put $in_\lambda(o)$ copies of o into T.
(2) Remove a bel a from L. For all bels b which are λ-adjacent from a, try to find one copy of b in T.
   a. If successful, remove this copy of b from T.
   b. If not, then put b into L and S and put $in_\lambda(b) - 1$ copies of b into T.
(3) Check if L is empty.
   a. If it is, STOP.
   b. If it is not, start again at Instruction (2).

Prior to proving the correctness of this algorithm, two remarks are in order. First, in Instruction (2)b, $in_\lambda(b) - 1$ is guaranteed to be nonnegative. This is because we only get to this point in the algorithm for a b which is λ-adjacent from a bel a and, so, $in_\lambda(b) \ge 1$. Second, the potential efficiency of this algorithm comes from the fact that we check for membership in T (rather than in S). While S keeps getting bigger and bigger as the algorithm is executed, due to Instruction (2)b, the same is not true for T: elements from T will be repeatedly removed due to Instruction (2)a. Hence the size of T is likely to be a small fraction of the size of S after the algorithm has been performing for a while on a large data set.

The essence of the proof of correctness is given in the next lemma. To state it easily we use a couple of abbreviations. We let $n_T(b)$ abbreviate "the number of copies of the bel b in the list T." The other definition is more complicated. The value of $n_a(b)$ is 0 if $b \notin S$ and is "$in_\lambda(b)$ less the number of bels in $S - L$ which are λ-adjacent to the bel b" otherwise.

Lemma 3.1. Both just prior and just after the execution of Instruction (2) in the General Bel-Tracking Algorithm it is the case that, for every bel b,

(i) $n_T(b) = n_a(b)$,

(ii) the bel b has so far been put into L and S — either due to Instruction (1) or due to Instruction (2)b — at most once, and

(iii) if the bel b is in S, then i is λ-connected in B to b.

Proof. Consider the situation just after the execution of Instruction (1). We have that $n_T(o) = in_\lambda(o) = n_a(o)$ and the bel o has so far been put into L and S exactly once. For any bel b other than o, $n_T(b) = 0 = n_a(b)$ and b has not so far been put into L and S even once. Since only o has been put into S, (iii) is clearly satisfied at this time.

Now we show that if (i), (ii) and (iii) hold just prior the execution of Instruction (2), then they also hold just after its execution. This is sufficient, since the situation cannot change as a result of Instruction (3). Assume therefore that we are just at the beginning of executing Instruction (2). This means that at this time L is not empty and so we remove a bel a from it. This a must have been put into L and S earlier on and, since nothing is ever removed from S, we must have that $a \in S$. We leave it to the reader to supply the easy proof of the fact that for those bels which are not λ-adjacent from a, the inductive step is valid. For the bels b which are λ-adjacent from a, we study separately two possibilities.

Case a: we find a copy of b in T. In this case, by Instruction (2)a, we remove this b from T. This reduces $n_T(b)$ by 1. Since b was in T, it also had to be in S (nothing is ever put into T without being put into S at the same time). Therefore the applicable part of the definition of $n_a(b)$ is that it is $in_\lambda(b)$ less the number of bels in $S - L$ which are λ-adjacent to the bel b. The only thing that changes in this definition is that a, which is λ-adjacent to b, got removed from L. Hence $n_a(b)$ is also decreased by 1, proving (i) of the lemma for the bel b. Since nothing is put into L and S, (ii) and (iii) are also valid at the end of executing Instruction (2).

Case b: there is no copy of b in T; i.e., $n_T(b) = 0$. First we show that under these circumstances, it cannot be the case that b has been previously put into L and S. This is so, since otherwise prior to the beginning of Instruction (2), the applicable part of the definition of $n_a(b)$ would be "$in_\lambda(b)$ less the number of bels in $S - L$ which are λ-adjacent to the bel b." Since at that time the bel a is still in L, the value of $n_a(b)$ has to be positive, contradicting the truth of (i) in the induction hypothesis. Upon executing Instruction (2)b, b has been put into L and S (for the first time) and $n_T(b) = in_\lambda(b) - 1$. There is at least one bel, namely a, which is in $S - L$ and is λ-adjacent to b. There cannot be another one, since whenever a bel is put into S, it is also put into L at the same time, and so if at a later time it is no longer in L, then it must have been removed from it. At that time, b would have been put into L and S and we would not be in Case b. This shows that in this case too, $n_T(b) = n_a(b)$ just after the execution of Instruction (2). Finally, since $a \in S$ just prior to the execution of Instruction (2), i is λ-connected in B to a, by (iii) of the induction hypothesis. The same must be true for the new bel b in S, since a is λ-adjacent to it. □

Theorem 3.2. If (V, π, f) is a binary picture, λ is a bel-adjacency and o is a bel, such that the λ-component of B which contains o is finite, then the General Bel-Tracking Algorithm terminates in a finite number of steps and, at that time, S is the λ-component of B which contains o.

Proof. From (iii) of the previous lemma it follows that anything that gets put into S is in the finite λ-component of B which contains i. Termination in a finite number of steps now easily follows from (ii) of the previous lemma: since each of the finitely many bels in the λ-component of B which contains i is put into L at most once (and nothing else ever gets put into L) and in each execution of Instruction (2) of the algorithm a bel is removed from L, sooner or later L has to become empty and the algorithm will stop due to Instruction (3). At that time, as all through the execution of the algorithm, i is λ-connected in B to every element of S. That the converse is also true (and hence S is the λ-component of B which contains i) can be shown as follows. For any bel b of the λ-component of B which contains i, there is a λ-path $\langle b^{(0)}, \cdots, b^{(K)} \rangle$ from i to b. It is a trivial matter to show by induction that, for $1 \le k \le K$, $b^{(k-1)}$ will get put into L and S and, since L gets eventually emptied, $b^{(k-1)}$ must get removed from L, resulting in $b^{(k)}$ being put into L and S (provided that it is not in T, which would imply that it has been put into L and S in some previous step). □

This theorem shows that the General Bel-Tracking Algorithm is powerful stuff. Its practical usefulness depends on two properties of the bel-adjacency λ. The first is, how easy is it to compute the bels λ-adjacent from a given bel? Clearly, the efficiency of executing Instruction (2) depends on this (as well as on how easy it is to determine for a bel whether or not it is in T). The other property has to do with the usefulness of the resulting boundaries: are they in fact of the form $\partial(O,Q)$ for some appropriately specified O and Q? In the following section we will positively answer these questions for some specific choices of the bel-adjacency in certain digital spaces.

4. BOUNDARY TRACKING ON HYPER-CUBES

We now apply the General Bel-Tracking Algorithm to the tracking of boundaries in the spaces (Z^N, ω_N) with $N \geq 2$. Let, for $1 \leq n \leq N$, u^n denote the *unit vector in direction n*, which is defined as the element of Z^N for which $u_n^n = 1$ and all other components are 0. A *basic digraph for* Z^N is a pair (M, ρ) for which:

(i) $M = \{ u^n \mid 1 \leq n \leq N \} \cup \{ -u^n \mid 1 \leq n \leq N \}$;

(ii) $\rho = \bigcup_{1 \leq i < j \leq N} \rho_{i,j}$, where $\rho_{i,j}$ is exactly one of

A. $\{(u^i, u^j), (u^j, -u^i), (-u^i, -u^j), (-u^j, u^i)\}$,

B. $\{(u^i, -u^j), (-u^j, -u^i), (-u^i, u^j), (u^j, u^i)\}$,

C. $\emptyset$.

Note that it is permitted in the definition of ρ that different options (A, B, or C) be selected for specifying $\rho_{i,j}$ for different pairs of i and j. We refer to elements of M as the *nodes* and elements of ρ as the *arcs* of the basic digraph.

As an example, consider Figure 4.1. In this case $N = 3$ and so the number of nodes in the basic digraph is six. Also, there are three sets of arcs $\rho_{i,j}$ (since $1 \leq i < j \leq 3$) and here we have selected Option A for all three. Selecting Option B would correspond to reversing the direction of the corresponding cycle of arrows (e.g., choosing Option B for $\rho_{2,3}$ would correspond to reversing the horizontal cycle of arrows on the vertical faces of the cube in Figure 4.1); such choices are arbitrary and, as will be easily seen from the material that follows, they make no difference to the nature of the tracked boundaries. (On the other hand, choosing Option C instead of Option A can result in a very different boundary being tracked by the General Bel-Tracking Algorithm.)

The purpose of introducing basic digraphs is that they can be used to define adjacencies between bels in binary pictures over (Z^N, ω_N). We now assume that we are given a basic digraph (M, ρ) and a binary picture (Z^N, ω_N, f) and explain how these are used to define an adjacency λ on the set of bels of (Z^N, ω_N, f). If $b = (c, d)$ is such a bel, then we must have that $f(c) = 1$, $f(d) = 0$, and $(c, d) \in \omega_N$. From the definition of ω_N it follows that $d - c$ is a node of the given basic digraph (M, ρ). Every node ρ-adjacent from $d - c$ gives rise to single bel λ-adjacent from b. Because of the utility of this in the proofs that follow, we in fact define λ as $\bigcup_{1 \leq i < j \leq N} \lambda_{i,j}$, where every node $\rho_{i,j}$-adjacent from $d - c$ gives rise to single bel $\lambda_{i,j}$-adjacent from b.

In following the details of the specification of how a node u which is $\rho_{i,j}$-adjacent from $d - c$ gives rise to a unique bel which is $\lambda_{i,j}$-adjacent from (c, d), consult Figure 4.2. (In that figure we are illustrating the situation when u has been selected using $\rho_{1,3}$. Three separate cases are considered: in (i) $d - c = -u^3$ and, hence, $u = u^1$, in (ii) $d - c = u^1$ and, hence, $u = u^3$, and in (iii) $d - c = u^3$ and, hence, $u = -u^1$.) We distinguish between three mutually exclusive possibilities. The first is that $d + u$ is a 1-spel (this corresponds to (iii) in Figure 4.2). In this case the bel $\lambda_{i,j}$-adjacent from (c, d) is specified to be $(d + u, d)$. The second is that $d + u$ is a 0-spel and $c + u$ is a 1-spel (this corresponds to (ii) in Figure 4.2). In this case the bel $\lambda_{i,j}$-adjacent from (c, d) is specified to be $(c + u, d + u)$. The third is that both $d + u$ and $c + u$ are 0-spels (this corresponds to (i) in Figure 4.2). In this case the bel $\lambda_{i,j}$-adjacent

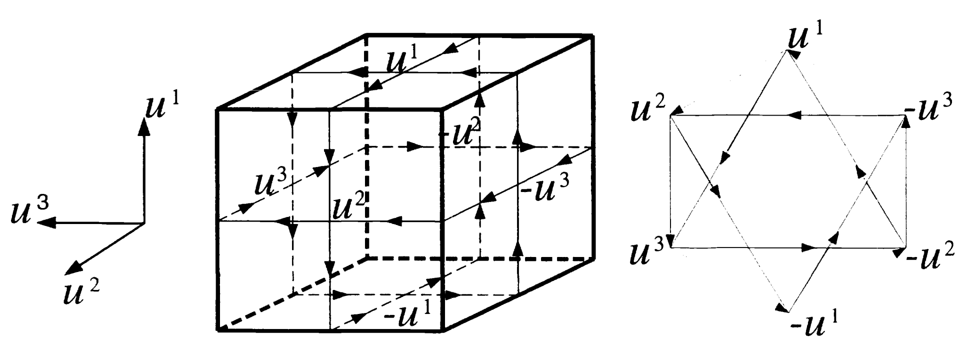

Figure 4.1. Illustration of a basic digraph (on the right) and its interpretation as a set of directions taken while tracking the boundary between a single spel and the set of all other spels in Z^3 (in the middle).

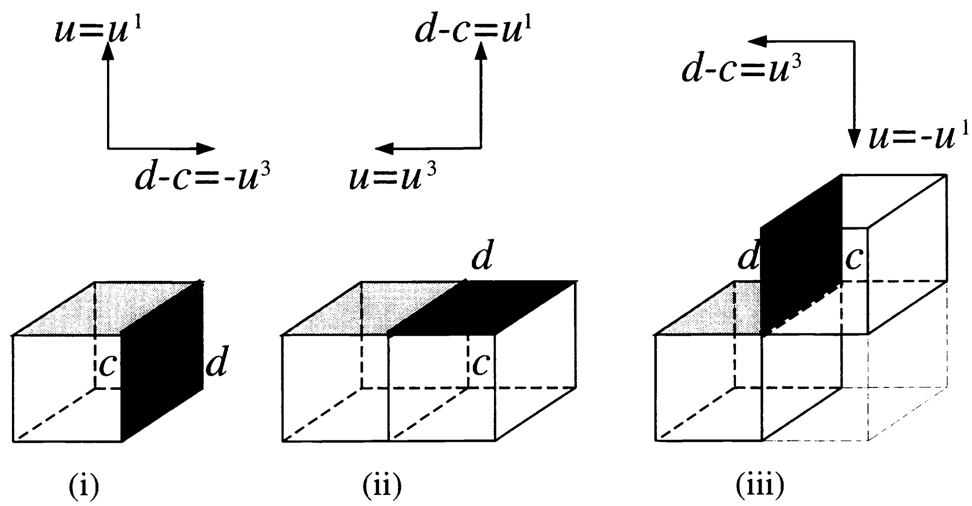

Figure 4.2. Illustration of the definition of bel-adjacency using the basic digraph of Figure 4.1. In each case, the darkly-shaded bel is $\lambda_{1,3}$-adjacent to the lightly-shaded one.

from (c, d) is specified to be $(c, c + u)$.

This definition provides a satisfactory reply to one of the questions raised in the last paragraph of the previous section: how easy is it to compute the bels λ-adjacent from a given bel? For the definition given in this section the answer is "it is very easy." A basic digraph is an easily represented object: for each of the $2N$ nodes, there are at most $N - 1$ nodes adjacent from it or to it. For any bel $b = (c, d)$, we have for some i $(1 \leq i \leq N)$ that $d - c \in \{u^i,\ -u^i\}$. For each (of the $N - 1$) $j \neq i$, there may (or may not) be a u which is $\rho_{i,j}$-adjacent from $d - c$. If there is such a u, it gives rise to a single bel $\lambda_{i,j}$-adjacent from (c, d) and the specification of this bel (as illustrated in Figure 4.2) requires the checking of the value of the binary function f at at most two spels.

The desirable behavior of the General Bel-Tracking Algorithm (as expressed by the theorem in the previous section) is guaranteed only if the adjacency λ is a bel-adjacency; i.e., if λ^* is symmetric. (The other condition in the definition of a bel-adjacency is automatically satisfied, since it immediately follows from the definition of λ that, for any bel b, $in_\lambda(b) \leq N - 1$.) Prior to discussing whether or not this is the case for the $\lambda = \bigcup_{1 \leq i < j \leq N} \lambda_{i,j}$ as defined in this section, we prove a technical lemma of a more general applicability.

Lemma 4.1. Let ρ be a binary relation on a set M and L be a finite subset of M.

(i) The transitive closure ρ^* of ρ is symmetric, provided that, for all $(c, d) \in \rho$, (d, c) is in ρ^*.

(ii) If every element of L has exactly one element of L ρ-adjacent from it and at least one element of L ρ-adjacent to it, then every element of L has exactly one element of L ρ-adjacent to it.

(iii) If every element of L has exactly one element of L ρ-adjacent from it and exactly one element of L ρ-adjacent to it, then for any c and d in L such that c is ρ-adjacent to d we have that d is ρ-connected in L to c.

Proof. Assume that, for all $(c, d) \in \rho$, (d, c) is in ρ^*. We prove by induction on K that if $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ is a ρ-path, then $(c^{(K)}, c^{(0)}) \in \rho^*$. This is clearly the case, by the definition of transitive closure, if $K = 0$. If $K > 0$, then $(c^{(K)}, c^{(K-1)}) \in \rho^*$ (by the assumption) and $(c^{(K-1)}, c^{(0)}) \in \rho^*$ (by the induction hypothesis). Since ρ^* is clearly transitive, we have that $(c^{(K)}, c^{(0)}) \in \rho^*$. This proves (i).

We now define the binary relation τ on L by $(c, d) \in \tau$ if, and only if, $c \in L$, $d \in L$, and $(c, d) \in \rho$. Since it is assumed in (ii) that every element of L has exactly one element of L ρ-adjacent from it, the number of elements in τ must be exactly the number of elements in L. This shows that there cannot be an element of L which has more than one element of L ρ-adjacent to it, since otherwise we would have more elements in τ than in L (since it is also assumed that every element of L has at least one element of L ρ-adjacent to it).

To complete the proof, let us assume that the premise of (iii) is satisfied and c and d in L are such that c is ρ-adjacent to d. We define an infinite sequence of elements of L as follows: $c^{(0)} = d$ and, for $k \geq 1$, $c^{(k)}$ is the unique element of L that is ρ-adjacent from $c^{(k-1)}$. Since L is finite, there must be a smallest positive integer l such that $c^{(l)} = c^{(m)}$, for some $m < l$. If m were positive, then the uniqueness of the element of L that is ρ-adjacent to $c^{(l)} = c^{(m)}$ would imply that $c^{(l-1)} = c^{(m-1)}$, which would contradict the minimality of l. Hence, $c^{(l)} = c^{(0)} = d$. Again by the uniqueness of the element of L which is ρ-adjacent to $c^{(l)} = d$, we must have that $c^{(l-1)} = c$, and so $\langle c^{(0)}, \cdots, c^{(l-1)} \rangle$ is a ρ-path in L from d to c. □

The rest of our results are proved under the assumption that the set of bels in (Z^N, ω_N, f) is finite. The assumption is not quite necessary, it could be replaced by the finiteness of the boundary that we intend to track. However, the additional generality is not worth the resulting additional awkwardness of proofs: in practice, we tend to deal with binary images in which the set of 1-spels is finite and, consequently, the set of bels is also finite.

Theorem 4.2. Let $(M, \bigcup_{1 \le i < j \le N} \rho_{i,j})$ be a basic digraph for Z^N. Assume that the set of all bels B in the binary picture (Z^N, ω_N, f) is finite and not empty and, for $1 \le i < j \le N$, define the binary relation $\lambda_{i,j}$ on B as above. Let $B_{i,j}$ be the subset of those bels $b = (c, d)$ for which $d - c \in \{u^i, -u^i, u^j, -u^j\}$. Then, for $1 \le i < j \le N$ such that $\rho_{i,j} \ne \emptyset$ and for every bel b in $B_{i,j}$, the following hold.

(i) There is one, and only one, bel b' that is $\lambda_{i,j}$-adjacent from b. Furthermore, $b' \in B_{i,j}$.

(ii) There is one, and only one, bel b' in $B_{i,j}$ that is $\lambda_{i,j}$-adjacent to b. Furthermore, b is $\lambda_{i,j}$-connected in $B_{i,j}$ to b'.

Moreover, $\lambda = \bigcup_{1 \le i < j \le N} \lambda_{i,j}$ is a bel-adjacency for (Z^N, ω_N, f).

Proof. By the definition of $\rho_{i,j}$, there is one and only one u which is $\rho_{i,j}$-adjacent from $d - c$ and this u is in $\{u^i, -u^i, u^j, -u^j\}$. According to the definition of $\lambda_{i,j}$, u gives rise to a unique bel $b' = (c', d')$ that is $\lambda_{i,j}$-adjacent from (c, d). Looking at the three possibilities in the definition of $\lambda_{i,j}$, we see that $d' - c'$ is (respectively) $d - (d + u) = -u$, $(d + u) - (c + u) = d - c$, or $(c + u) - c = u$. In any case, $b' \in B_{i,j}$. This proves (i).

In order to prove (ii) we first distinguish between three mutually exclusive possibilities (corresponding to the three cases represented in Figure 4.3, with $b = (c, d)$ corresponding to the lightly-shaded bel in that figure) and show that in each one of them there is at least one bel $b' = (c', d')$ in $B_{i,j}$ that is $\lambda_{i,j}$-adjacent to b. We again denote by u the element of $\{u^i, -u^i, u^j, -u^j\}$ which is $\rho_{i,j}$-adjacent from $d - c$.

The first possibility is that $d - u$ is a 1-spel (this corresponds to (iii) in Figure 4.3). In this case we define $c' = d - u$ and $d' = d$. Then $b' = (c', d')$ is a bel, $d' - c' = u$, and so $b' \in B_{i,j}$. From the definition of $\rho_{i,j}$ it follows (since u is $\rho_{i,j}$-adjacent from $d - c$) that $u' = c - d$ is $\rho_{i,j}$-adjacent from $d' - c' = u$. Since $d' + u' = d + (c - d) = c$ is a 1-spel, the definition of $\lambda_{i,j}$ implies $b' = (c', d')$ is $\lambda_{i,j}$-adjacent to $(d' + u', d') = (c, d) = b$.

The second possibility is that $d - u$ is a 0-spel and $c - u$ is a 1-spel (this corresponds to (ii) in Figure 4.3). In this case we define $c' = c - u$ and $d' = d - u$. Then $b' = (c', d')$ is a bel, $d' - c' = d - c$, and so $b' \in B_{i,j}$. Since u is $\rho_{i,j}$-adjacent from $d - c = d' - c'$ and $d' + u = d$ is a 0-spel and $c' + u = c$ is a 1-spel, the definition of $\lambda_{i,j}$ implies

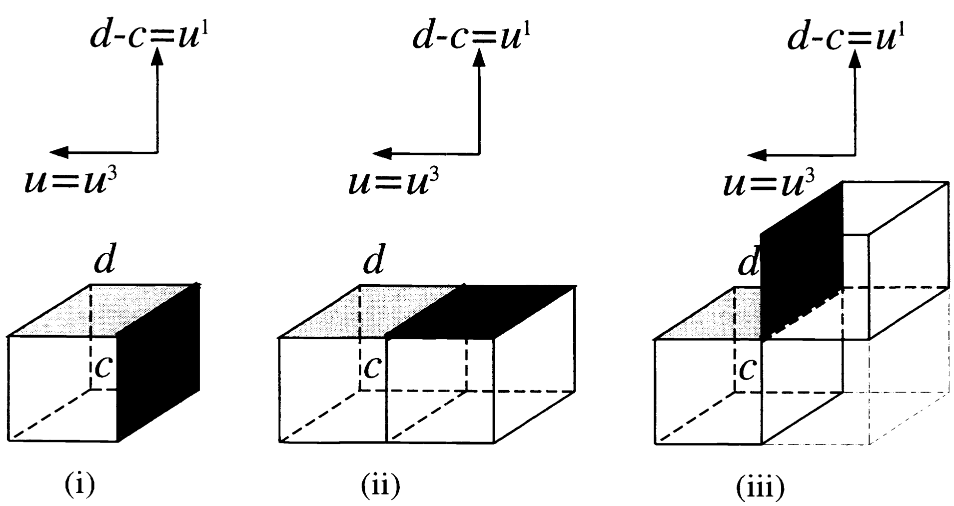

Figure 4.3. Illustration of Theorem 4.2(ii) using the basic digraph of Figure 4.1. In each case, the darkly-shaded bel is $\lambda_{1,3}$-adjacent to the lightly-shaded one.

$b' = (c', d')$ is $\lambda_{i,j}$-adjacent to $(c' + u, d' + u) = (c, d) = b$.

The third possibility is that both $d - u$ and $c - u$ are 0-spels (this corresponds to (i) in Figure 4.3). In this case we define $c' = c$ and $d' = c - u$. Then $b' = (c', d')$ is a bel, $d' - c' = -u$, and so $b' \in B_{i,j}$. From the definition of $\rho_{i,j}$ it follows (since u is $\rho_{i,j}$-adjacent from $d - c$) that $u' = d - c$ is $\rho_{i,j}$-adjacent from $d' - c' = -u$. Since $d' + u' = (c - u) + (d - c) = d - u$ and $c' + u' = c + (d - c) = d$ are both 0-spels, the definition of $\lambda_{i,j}$ implies $b' = (c', d')$ is $\lambda_{i,j}$-adjacent to $(c', c' + u') = (c, d) = b$.

This shows that, for any bel b, there is at least one bel b' in $B_{i,j}$ that is $\lambda_{i,j}$-adjacent to b. According to (i), it is also the case that there is exactly one bel b' in $B_{i,j}$ that is $\lambda_{i,j}$-adjacent from b. Applying (ii) of the previous lemma we get that there is one, and only one, bel b' in $B_{i,j}$ that is $\lambda_{i,j}$-adjacent to b. Hence we can also apply (iii) of the previous lemma and obtain that b is $\lambda_{i,j}$-connected in $B_{i,j}$ to b'. This completes the proof of (ii).

To complete the proof, we make use of (i) of the previous lemma. If $(a, b) \in \lambda$, then $(a, b) \in \lambda_{i,j}$, for some $1 \le i < j \le N$ such that $\rho_{i,j} \neq \emptyset$. By (ii), this implies that b is $\lambda_{i,j}$-connected in $B_{i,j}$ to a and, consequently, $(b, a) \in \lambda^*$. Hence, λ^* is symmetric. □

Corollary 4.3. For $N \ge 2$, let (M, ρ) be a basic digraph for Z^N. Let (Z^N, ω_N, f) be a binary picture with a finite set B of bels and define the binary relation λ on B as above. Then, for any bel o, the General Bel-Tracking Algorithm terminates in a finite number of steps and, at that time, S is that λ-component of the set of bels which contains o.

We call a loop $\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)}, c^{(0)} \rangle$ a *unit square* if both $c^{(0)} \neq c^{(2)}$ and $c^{(1)} \neq c^{(3)}$. Consider a digital space (Z^N, ω_N) with N a positive integer. In such a space we call a loop $\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)}, c^{(0)} \rangle$ of length four a *unit lattice square* if the following conditions are satisfied. There exist integers i, j $(1 \le i \neq j \le N)$ and u, v $(|u| = |v| = 1)$ such that

$$c_i^{(1)} = c_i^{(0)} + u, \quad c_j^{(2)} = c_j^{(1)} + v, \quad c_i^{(3)} = c_i^{(2)} - u, \quad c_j^{(0)} = c_j^{(3)} - v \; . \tag{4.1}$$

We refer to the pairs $(c^{(0)}, c^{(2)})$ and $(c^{(1)}, c^{(3)})$ as the *diagonals* of the unit lattice square.

Lemma 4.4. In the digital space (Z^N, ω_N) with $N \ge 1$ a loop is a unit square if, and only if, it is a unit lattice square.

Proof. It is clear from (4.1) that a unit lattice square is a unit square. Conversely, let the loop $\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)}, c^{(0)} \rangle$ be a unit square in (Z^N, ω_N). Since the successive spels are proto-adjacent, they must differ by ± 1 in exactly one coordinate. Since $c^{(0)} \neq c^{(2)}$ and $c^{(1)} \neq c^{(3)}$, these coordinates cannot be the same for two consecutive pairs, and so must be the same (and with opposite signs of the difference) for the alternating pairs. □

Theorem 4.5. For $N \ge 2$, let $(M, \bigcup_{1 \le i < j \le N} \rho_{i,j})$ be a basic digraph for Z^N such that $\rho_{i,j} \neq \emptyset$ for any $1 \le i < j \le N$. Assume that the set of all bels B in the binary picture (Z^N, ω_N, f) is finite and define the binary relation λ on B as above. It is then the case that, for any bel $o = (c, d)$, the General Bel-Tracking Algorithm terminates in a finite number of steps and, at that time, $S = \partial(O, Q)$, where O is the δ_N-component of the set of 1-spels which contains c and Q is the ω_N-component of the set of 0-spels which contains d.

Proof. By the previous corollary we know that the algorithm terminates in a finite number of steps and that at that time S is that λ-component of B which contains o.

First we show inductively that both just prior and just after the execution of Instruction (2) in the General Bel-Tracking Algorithm $L \subset \partial(O, Q)$ and $S \subset \partial(O, Q)$. This is true just prior to the first execution of Instruction (2), since only o gets put into L and S due to Instruction (1) and, clearly, $o = (c, d) \in \partial(O, Q)$. Assume now the validity of the induction hypothesis just prior to some execution of Instruction (2). If a bel b gets put into L and S during the execution of Instruction (2), then there must have been an a removed from $L \subset \partial(O, Q)$ such that b is $\lambda_{i,j}$-adjacent from a. By considering the three possible ways that a can be $\lambda_{i,j}$-adjacent to b (look at Figure 4.2) we see that $a \in \partial(O, Q)$ implies that $b \in \partial(O, Q)$. (Essential use is made here of the fact that O is a δ_N-component of the set of 1-spels.) Since nothing is put into L or S during the execution Instruction (3), this completes the induction.

Next we show that at the termination of the General Bel-Tracking Algorithm S is 2-locally-Jordan. By definition, we need to show that $p_S\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)} \rangle$ is odd for any ω_N-path such that $(c^{(0)}, c^{(3)}) \in S$. (We have made use here of the fact that there is no ω_N-path of length 2 from a spel to a spel ω_N-adjacent from it.) The sought-after result is trivially true if $c^{(0)} = c^{(2)}$ or $c^{(1)} = c^{(3)}$. When neither of these is the case, then the loop $\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)}, c^{(0)} \rangle$ is (by definition) a unit square and, hence by Lemma 4.4 a unit lattice square. Looking at (4.1), we see that this implies that, for some $1 \le i < j \le N$, $c^{(3)} - c^{(0)} = c^{(2)} - c^{(1)} \in \{u^i, -u^i, u^j, -u^j\}$ and $c^{(3)} - c^{(2)} = c^{(0)} - c^{(1)} \in \{u^i, -u^i, u^j, -u^j\}$; see Figure 4.4. In that figure we distinguish between three cases, corresponding to the three cases of Figure 4.2.

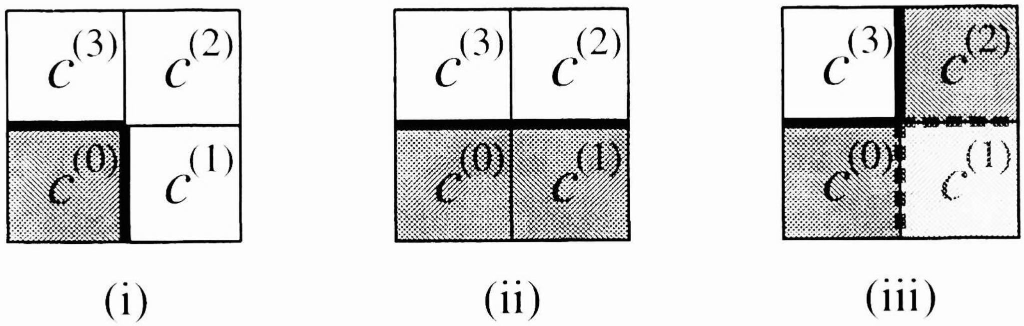

Figure 4.4. Illustration of the three cases in the proof which shows that the output of the General Bel-Tracking Algorithm based on the basic digraph $(M, \bigcup_{1 \le i < j \le N} \rho_{i,j})$ with $\rho_{i,j} \neq \emptyset$, for any $1 \le i < j \le N$, is 2-locally-Jordan. The spels shaded dark grey are 1-spels, the spel shaded light grey may be either a 1-spel or a 0-spel, all other spels are 0-spels.

In each case, the bel $(c^{(0)}, c^{(3)})$ in Figure 4.4 corresponds to the bel painted light grey in Figure 4.2. In Case (i), $f(c^{(1)}) = f(c^{(2)}) = 0$. Both the bels $(c^{(0)}, c^{(3)})$ and $(c^{(0)}, c^{(1)})$ are in $B_{i,j}$ (as defined in the statement of the previous theorem) and, since $\rho_{i,j} \neq \emptyset$, one of them has to be $\lambda_{i,j}$-adjacent to the other. The previous theorem implies that they are in the same λ-component of B. Since $(c^{(0)}, c^{(3)}) \in S$, this implies that $(c^{(0)}, c^{(1)}) \in S$. On the other hand, $(c^{(1)}, c^{(2)})$ and $(c^{(2)}, c^{(3)})$ are not bels and so cannot be in S. This proves that $ps\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)} \rangle = 1$ in Case (i). The proof in Case (ii) is very similar and is therefore not given. In Case (iii) there is an additional complication. One can argue, as in Case (i), to show that in Case (iii) $(c^{(2)}, c^{(3)}) \in S$, but the possibilities of $(c^{(0)}, c^{(1)}) \in S$ or $(c^{(1)}, c^{(2)}) \in S$ need also be investigated. Neither of these is true if $c^{(1)}$ is a 1-spel. On the other hand, if $c^{(1)}$ is a 0-spel, then both $(c^{(0)}, c^{(1)})$ and $(c^{(1)}, c^{(2)})$ are in $B_{i,j}$ and, consequently, either both are in S or neither is in S. In any case, we again have that $ps\langle c^{(0)}, c^{(1)}, c^{(2)}, c^{(3)} \rangle$ is odd.

Next we show that at the termination of the General Bel-Tracking Algorithm S is near-Jordan. We use Lemma 2.4. We have just shown that S is a 2-locally-Jordan surface in the digital space (Z^N, ω_N), which is 2-simply connected, by Theorem 2.5. Since S is a set of bels, it has to be antisymmetric. To show that it is near-Jordan it is sufficient to show (according to Lemma 2.4) that, for any two elements in the immediate exterior of S, there is an ω_N-path between them which does not cross S. Suppose that g and h are in $IE(S)$. That implies that there exist spels g' and h', such that both (g', g) and (h', h) are in S. Since S is a λ-component of B, there is λ-path in S from (g', g) to (h', h). Since a λ-path is a sequence of consecutively λ-adjacent bels, the task of this paragraph is done if we can show that whenever a bel (g', g) is λ-adjacent to a bel (h', h), then there is an ω_N-path from g to h that does not cross B. This follows, since (g', g) must be $\lambda_{i,j}$-adjacent to (h', h), for some $1 \le i < j \le N$, and in each of the three cases in the definition of $\lambda_{i,j}$ the required result follows trivially. (For example, the three relevant ω_3-paths in Figure 4.2 are: (i) $\langle d, d+u, c+u \rangle$, (ii) $\langle d, d+u \rangle$, and (iii) $\langle d \rangle$.)

Now we complete the proof by showing that at the termination of the General Bel-Tracking Algorithm $S = \partial(O, Q)$. We already know that $S \subset \partial(O, Q)$. To prove the converse, consider any $(g, h) \in \partial(O, Q)$; i.e., $g \in O$, $h \in Q$, $(g, h) \in \omega_N$. Since $g \in O$, there is a δ_N-path in O from c to g from which we can create an ω_N-path $\langle c^{(0)}, \cdots, c^{(K)} \rangle$ from c to g with the following property: for $0 \le k \le K$, either $f(c^{(k)}) = 1$ or $f(c^{(k)}) = 0$ and $f(c^{(k-1)}) = f(c^{(k+1)}) = 1$. In the latter case, $(c^{(k-1)}, c^{(k)})$ and $(c^{(k+1)}, c^{(k)})$ are bels such that one of them is λ-adjacent to the other and so either both are in S or neither is in S. It follows therefore that $ps\langle c^{(0)}, \cdots, c^{(K)} \rangle = 0$. Since $h \in Q$, there is an ω_N-path $\langle d^{(0)}, \cdots, d^{(L)} \rangle$ in Q from h to d; clearly $ps\langle d^{(0)}, \cdots, d^{(L)} \rangle = 0$. It follows that $(g, h) \in S$, for otherwise the crossing parity through the near-Jordan surface S would be even for the ω_N-path $\langle c = c^{(0)}, \cdots, c^{(K)} = g, h = d^{(0)}, \cdots, d^{(L)} = d \rangle$, contradicting Theorem 2.1. □

ACKNOWLEDGMENT

This work is supported by the National Science Foundation Grant DMS-9612077.

REFERENCE

[1] G.T. Herman. *Geometry of Digital Spaces*. Birkhauser Boston, Cambridge, MA, 1998.