SOLUTIONS TO AFFINE GENERALIZED EQUATIONS USING PROXIMAL MAPPINGS

STEPHEN C. BILLUPS AND MICHAEL C. FERRIS

The normal map has proven to be a powerful tool for solving generalized equations of the form: find $z \in C$, with $0 \in F(z) + N_C(z)$, where C is a convex set and $N_C(z)$ is the normal cone to C at z. In this paper, we use the T-map, a generalization of the normal map, to solve equations of the more general form: find $z \in \mathrm{dom}(T)$, with $0 \in F(z) + T(z)$, where T is a maximal monotone multifunction. We present a path-following algorithm that determines zeros of coherently oriented piecewise-affine functions, and we use this algorithm, together with the T-map, to solve the generalized equation for affine, coherently oriented functions F, and polyhedral multifunctions T. The path-following algorithm we develop here extends the piecewise-linear homotopy framework of Eaves to the case where a representation of a subdivided manifold is unknown.

1. Introduction. This paper is concerned with solving generalized equations (Robinson 1979b, Robinson 1983) of the form

(1) $$0 \in F(x) + T(x),$$

where T is a maximal monotone multifunction from $\mathbf{R}^n$ into $\mathbf{R}^n$ and F is a continuously differentiable function from an open set $\Omega \supset \mathrm{dom}(T)$ into $\mathbf{R}^n$. We recall that a monotone multifunction T is a point to set mapping such that for each (x^1, y^1), (x^2, y^2) in the graph of T,

$$\langle x^1 - x^2, y^1 - y^2 \rangle \geq 0,$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product, T is maximal if its graph is not properly contained in that of any other monotone multifunction and $\mathrm{dom}(T)$ represents the effective domain of T.

To date, most of the algorithmic development for generalized equations has been focused on the special case where $T := N_C$, the normal cone to a convex set C, defined by

$$T(z) = N_C(z) := \begin{cases} \{w \mid \langle w, y - z \rangle \leq 0, \ \forall y \in C\} & z \in C, \\ \varnothing & z \notin C. \end{cases}$$

It is well known that N_C is a maximal monotone multifunction (since, for example, it is the subdifferential of the indicator function of C; see Aubin and Ekeland 1984, Remark, page 194). This case, therefore, yields the generalized equation

(2) $$0 \in F(x) + N_C(x).$$

Many problems from mathematical programming, equilibrium, complementarity and other fields can be expressed in this form. For example, if $F := \nabla f$, then (2) represents the first order necessary optimality conditions for the problem

$$\text{minimize} \quad f(x)$$
$$\text{subject to} \quad x \in C.$$

Another important instance of (2) is the variational inequality problem, which is to find $z \in C$ such that

$$\langle y - z, F(z) \rangle \geq 0, \qquad \forall\, y \in C.$$

This problem is known to be equivalent to (2), see Robinson (1983).

As a final example, we mention the complementarity problem; this has a wealth of applications and appears in a variety of forms (Ferris and Pang 1997). The standard nonlinear complementarity problem is to find $x \in \mathbf{R}_+^n$ such that

$$F(x) \geq 0 \quad \text{and} \quad \langle x, F(x) \rangle = 0.$$

It is well known (Ferris and Pang 1997, Karamardian 1971) that complementarity problems can be reformulated as variational inequalities, and therefore can be treated in the context of the generalized equation (2).

A prominent tool for solving (2) is the *normal map* (Eaves 1971, Robinson 1992, Todd 1976). The normal map F_C for a function $F : \Omega \to \mathbf{R}^n$ and a nonempty, closed, convex set $C \supset \Omega$ is defined by

$$(3) \qquad\qquad F_C(x) := F(\pi_C(x)) + x - \pi_C(x),$$

where $\pi_C(x)$ denotes the projection (with respect to the Euclidean norm) of x onto the set C. The importance of the normal map lies in the fact that solving (2) is equivalent to finding a zero of the normal map F_C. Specifically, if x is a zero of the normal map, then $z := \pi_C(x)$ solves (2). Conversely, if z is a solution to (2), then $x := z - F(z)$ is a zero of the normal map. Thus, the problem of solving the generalized equation (2), which is expressed with set-valued functions, is reduced to finding a zero of a piecewise smooth, single-valued function.

Several algorithms have been developed based on this idea. Most of these algorithms are based on the theory of piecewise linear homotopies given by Eaves (1976). The specialization of the general algorithm given in Eaves (1976) determines a zero of the normal map by tracing the zero curve of a piecewise linear homotopy mapping. It is shown in Eaves and Lemke (1981) that Lemke's famous method (Lemke and Howson 1964) for solving linear complementarity problems is conceptually equivalent to path following in a corresponding piecewise linear system. This idea is easily extended to more general algorithms (Eaves 1978a, Eaves 1978b). Based on Eaves' work, Cao and Ferris (1996b) analyzed an algorithm for solving affine (i.e. $F(x)$ is affine) variational inequality problems over polyhedral sets and showed that Lemke's method processed matrices in the class P_C (Cao and Ferris 1996a).

For non-affine problems with C rectangular, Ralph (1994) proposed a Newton-based algorithm where at each iteration the Newton point is calculated by solving an affine variational inequality (AVI) that is a linearization of the normal map equation at the current point. This approach was developed computationally by Dirkse and Ferris (1995) to produce PATH, an efficient and robust code for solving mixed complementarity

problems, see also (Dirkse 1994, Dirkse and Ferris 1996). A similar approach was developed by Rutherford (1993).

In this paper, we begin to generalize the above class of algorithms by removing the restriction that T be the normal cone to a convex set. As a first step, this paper focuses on solving the *affine generalized equation* where F is an affine function and T is polyhedral (that is, the graph of T is the union of finitely many convex polyhedrons).

Our strategy is, in essence, a generalization of the AVI algorithm of Cao and Ferris. We use a generalization of the normal map to reformulate (1) as a zero finding problem of a piecewise affine function. To solve this zero finding problem, we present a path following algorithm for determining zeros of piecewise-affine functions.

Our path following algorithm is based on the piecewise-linear homotopy algorithm of Eaves (1976). However it is more flexible in the following sense. In Eaves' framework, the function is defined in terms of a subdivided manifold (see Definition 2.6) such that the function is affine when restricted to any single cell of this manifold. To implement the algorithm, we need an explicit representation of this subdivided manifold. In our application, such a representation is not readily available. Instead, the function is defined in terms of a more general collection of cells which do not necessarily comprise a subdivided manifold. We therefore generalize Eaves' framework to develop an algorithm for finding zeros of piecewise-affine functions that does not rely on a representation of a subdivided manifold. We describe this algorithm in detail in §2 and prove that under the assumption of coherent orientation, the algorithm finds a zero after a finite number of steps.

The remainder of the paper describes how the algorithm is used to solve affine generalized equations. In §3 we discuss a generalization of the normal map called the T-map. This mapping, which to our knowledge was introduced by Minty (1962), provides a means of reformulating generalized equations involving operators T that do not necessarily correspond to the normal cone of any set. After describing this reformulation, we focus on the case where T is polyhedral. This case can be viewed as an extension of the special case where $T := N_C$, with C a convex polyhedron. In §4 we describe how the algorithm presented in Section 2 can be used to solve affine generalized equations. In §5 we further specialize to the case where T is separable. Finally, in §6 we discuss how the algorithm can be applied to solve piecewise linear-quadratic programming problems (Sun 1986).

Some words about notation are needed. The notation $T : \mathbf{R}^n \rightrightarrows \mathbf{R}^n$ indicates that T is a point to set mapping, or *multifunction*, which maps points in $\mathbf{R}^n$ to subsets of $\mathbf{R}^n$. In discussing multifunctions, we may refer either to a point to set mapping, $T(\cdot)$, or to the graph of that mapping, T, which is the set $T := \{(x, y) | y \in T(x)\}$. The expression $T^{-1}(\cdot)$ is defined as a set inverse; i.e., $T^{-1}(y) := \{x | (x, y) \in T\}$. Further, $T^{-1} = \{(y, x) | (x, y) \in T\}$. The effective domain of T, is defined by $\mathrm{dom}(T) := \{x | T(x) \neq \varnothing\}$.

In discussing matrices and vectors, subscripts are used to refer to components. For example $A_i \cdot$, $A \cdot_j$, A_{ij} refer to the ith row, jth column, and (i, j)th entry of A, respectively. We may also use index sets to refer to specific components. For example if $\alpha = \{1, 3, 5\}$, then $x_\alpha = \{x_1; x_3; x_5\}$. Further, we use the MATLAB notation of a comma to separate columns, and a semicolon to separate rows. For example (x, y) is a row vector, whereas $(x; y)$ is a column vector. Unless otherwise indicated, all vectors are taken to be column vectors. Superscripts are used to indicate an iteration count, index, or some other label for matrices and vectors. In contrast, for scalars, sets, and functions, subscripts are used as labels.

For a set C, $\mathrm{aff}(C)$, $\mathrm{int}(C)$, $\mathrm{ri}(C)$, $\mathrm{rec}(C)$, $\mathrm{dim}(C)$, and $\delta(\cdot | C)$ refer to the affine hull, interior, relative interior, recession cone, dimension, and indicator function of C respectively. (see Rockafellar 1970 for definitions of these terms).

2. Algorithm for finding zeros of $\mathcal{M}$-PA maps. Our first task is to describe an algorithm for finding zeros of piecewise affine functions. The algorithm we present can be viewed as a generalization of Lemke's method for the linear complementarity problem and

of the algorithm for solving affine variational inequalities over convex polyhedral sets that was described in Cao and Ferris (1996b). The theoretical basis for the algorithm is derived from the theory of piecewise-linear homotopies given in Eaves (1976). However, we emphasize that the theory described in Eaves (1976) assumes a representation of a certain subdivided manifold. One of our major tasks in this section is to remove this assumption so that the algorithm can be applied in a more general setting where such a subdivided manifold is not readily available. This extension is necessary for application of the algorithm to the T-map as described in §4.

In order to describe the algorithm carefully, we need some preliminary definitions.

DEFINITION 2.1 (cell).   A polyhedral convex set $\sigma \subset \mathbf{R}^n$ is called a cell. If $\dim(\sigma) = k$ then σ is called a k-cell. Let $\sigma := \{x \mid Ax \leq a\}$, where $A \in \mathbf{R}^{p \times n}$, and $a \in \mathbf{R}^p$, with p a nonnegative integer. Then (p, A, a) is said to represent σ. If p is the smallest number for which a representation of σ exists, then (p, A, a) is called a minimal representation of σ. A set $\tau \in \mathbf{R}^n$ is called a face of σ if for some set of indices $\alpha \subset \{1, \ldots, p\}$, $\tau = \{x \in \sigma : A_\alpha . x = a_\alpha\}$. If $\dim(\tau) = i$, then τ is called an i-face of σ.

Clearly any cell has a minimal representation.

DEFINITION 2.2 (piecewise affine).   Let $\mathcal{M}$ be a finite collection of n-cells and let $M := \cup_{\sigma \in \mathcal{M}} \sigma$. A function $F : M \to \mathbf{R}^m$ is said to be piecewise-affine with respect to $\mathcal{M}$, denoted $\mathcal{M}$-PA, if for each $\sigma \in \mathcal{M}$, $F_{|\sigma}$ (i.e. the restriction of F to σ) is affine. If F is $\mathcal{M}$-PA for some $\mathcal{M}$ satisfying the above assumptions, then we say that F is piecewise affine.

Note that in the above definition, if M is connected, then the function F must be continuous on M, because F must be single-valued on the boundaries between cells. Furthermore, in contrast to the work of Eaves (1976), $\mathcal{M}$ is not required to correspond to a subdivided manifold. An alternative definition and development of piecewise affine functions is given by Scholtes (1994) for the case where F is a continuous function on $\mathbf{R}^n$.

DEFINITION 2.3 (function representation).   Let $\mathcal{M}$ be a collection of n-cells in $\mathbf{R}^n$, let F be a $\mathcal{M}$-PA function, and let σ be an n-cell of $\mathcal{M}$. Let $b^\sigma \in \mathbf{R}^m$ and let B^σ be an $m \times n$ matrix. (B^σ, b^σ) is said to represent F on σ if $F(x) = B^\sigma x + b^\sigma$ for all $x \in \sigma$.

We now describe an algorithm to find a zero of an $\mathcal{M}$-PA function G, for a given finite collection of n-cells $\mathcal{M}$ whose union is $\mathbf{R}^n$. We will assume that representations of the cells of $\mathcal{M}$ and of the map G have already been constructed. The basis of the algorithm is to construct a piecewise affine homotopy mapping $F(x, \mu)$ with the following properties

1. $(x^*, 0)$ is a zero of F if and only if x^* is a zero of G.

2. A point (x^1, μ_1), and a direction $(d^1, -1)$ is known such that $\mu_1 \geq 0$ and $F(x^1 - \mu d^1, \mu^1 + \mu) = 0$ for all $\mu \geq 0$.

The algorithm uses a method described by Eaves (1976) to trace the zero curve of F, proceeding in the direction $(d^1, -1)$ from the starting point (x^1, μ_1). To prove that the algorithm finds a solution in a finite number of steps, we restrict ourselves to the case where G is *coherently oriented*:

DEFINITION 2.4 (coherent orientation).   Let G be an $\mathcal{M}$-PA map with representation (B^σ, b^σ) on each $\sigma \in \mathcal{M}$. We say that G is coherently oriented if

$$\operatorname{sgn}(\det(B^\sigma))$$

is nonzero and constant for all σ in $\mathcal{M}$, where

$$\operatorname{sgn}(x) := \begin{cases} -1 & x < 0 \\ 0 & x = 0, \\ 1 & x > 0. \end{cases}$$

Since $\mathcal{M}$ is finite and $\cup_{\sigma \in \mathcal{M}} \sigma = \mathbf{R}^n$, it follows that $\mathbf{R}^n = \cup_{\sigma \in \mathcal{M}} \operatorname{rec}(\sigma)$, and further that there is a σ such that $\operatorname{int}(\operatorname{rec}(\sigma)) \neq \varnothing$. Choose d such that $-d \in \operatorname{int}(\operatorname{rec}(\sigma))$. Then for any x^0 in $\mathbf{R}^n$, and for all μ sufficiently large, $x^0 - \mu d \in \operatorname{int}(\sigma)$.

In the AVI algorithm described by Cao and Ferris, the cell σ and the direction d were constructed by finding an extreme point x^e of the set C. The cell was then given by $\sigma := x^e + N_C(x^e)$, and the direction d was chosen such that $-d$ was in the interior of $N_C(x^e)$. For our algorithm, rather than constructing the cell and direction, we can rely instead on the fact that since $\mathbf{R}^n = \cup_{\sigma \in \mathcal{M}} \operatorname{rec}(\sigma)$, then for any direction d, there will be a cell σ_d for which $-d \in \operatorname{rec}(\sigma_d)$. Note further that for each cell σ, the boundary $\operatorname{rec}(\sigma)\backslash\operatorname{int}(\operatorname{rec}(\sigma))$ of $\operatorname{rec}(\sigma)$ has Lebesgue measure zero. Therefore, since the number of cells is finite, $\cup_{\sigma \in \mathcal{M}} \operatorname{rec}(\sigma)\backslash\operatorname{int}(\operatorname{rec}(\sigma))$ has measure zero. Thus, for almost all d, there will be a cell σ_d for which $-d \in \operatorname{int}(\operatorname{rec}(\sigma_d))$.

If x^0 is any point in $\mathbf{R}^n$, then for all μ sufficiently large, $x^0 - \mu d$ will lie interior to the cell σ_d. In other words, the cell can be chosen simply by picking an arbitrary d and proceeding in the direction $-d$ until a cell σ_d is reached for which $-d$ is in the recession cone of σ_d. For almost all d (excepting a set of Lebesgue measure zero), $-d$ will be in the *interior* of $\operatorname{rec}(\sigma_d)$. We note, however, that for some special cases, construction of an extreme point may still be preferable.

Once d and σ_d have been identified, the homotopy map can be constructed. Let (B, b) be the representation of G in σ_d. Define a function $F : \mathbf{R}^n \times \mathbf{R}_+ \to \mathbf{R}^n$ by

$$(4) \qquad F(x, \mu) := G(x) + \mu B d.$$

Note that $F(x, 0) = 0$ exactly when $G(x) = 0$. Under the assumption that G is coherently oriented, B is invertible. Let $x^0 := -B^{-1}b$ and define

$$w(\mu) = x^0 - \mu d.$$

Then, since $-d \in \operatorname{int}(\operatorname{rec}(\sigma_d))$, there exists $\sigma_0 \geq 0$ such that $w(\mu) \in \operatorname{int}(\sigma_d)$, $\forall \mu > \mu_0$. Thus, for $\mu \geq \mu_0$,

$$
\begin{aligned}
F(w(\mu), \mu) &= G(w(\mu)) + \mu B d \\
&= B w(\mu) + b + \mu B d \\
(5) \qquad &= B(x^0 - \mu d) + b + \mu B d \\
&= -b - \mu B d + b + \mu B d \\
&= 0.
\end{aligned}
$$

By choosing $\mu_1 > \mu_0$, $x^1 = w(\mu_1)$, and $d^1 = d$, we see that F satisfies the conditions needed for the homotopy map. We are now ready to state the algorithm, which is given in Figure 1. Note that by normalizing d in the discussion above to be a unit vector, we can start the algorithm from the point (x^1, μ_1) constructed above with $\sigma_1 := \sigma_d$.

Some comments about Algorithm AGE are in order:

1. Most of the work in the algorithm is in step 8 where the direction (d^{k+1}, v_{k+1}) is calculated. At the end of this section, in Theorem 2.13, we show that $B^{k+1} - B^k$ is a rank-1 matrix. Thus, an efficient implementation of the algorithm can be obtained by keeping the matrices B^k in factored form and performing rank-1 updates of the factors at each step of the algorithm.

Given a finite collection of n-cells $\mathcal{M}$ such that $\cup_{\sigma \in \mathcal{M}} \sigma = \mathbf{R}^n$, and an $\mathcal{M}$-PA function G on $\mathbf{R}^n$. Let G have representation (B^k, b^k) on $\sigma_k \in \mathcal{M}$.

1) [**Initialization**] Determine $(x^1, \mu_1, \sigma_1, d^1)$ satisfying
$$\|d^1\| = 1,$$
$$x^1 \in \mathrm{int}(\sigma_1),$$
$$x^1 - \mu d^1 \in \mathrm{int}(\sigma_1), \ \forall \mu \geq 0,$$
$$B^1 x^1 + \mu_1 B^1 d^1 + b^1 = 0.$$

2) $v_1 := -1$.

Repeat for $k = 1, 2, \ldots$

3) $\theta_k := \sup\{\theta | x^k + \theta d^k \in \sigma_k, \ \mu_k + \theta v_k \geq 0\}$.

    If $\theta_k = +\infty$, then

        4) **output**("ray termination"); **return.**

    Else

        5) $x^{k+1} := x^k + \theta_k d^k$

        6) $\mu_{k+1} := \mu_k + \theta_k v_k$

        If $\mu_{k+1} = 0$ then

        7) **output**("solution found at", x^{k+1}); **return.**

        Else

        8) determine σ_{k+1} (possibly using lexicographic ordering), d^{k+1}, and v_{k+1} such that
$$x^{k+1} \in \sigma_{k+1},$$
$$B^{k+1} d^{k+1} + v_{k+1} B^1 d^1 = 0,$$
$$\|d^{k+1}\| = 1,$$
$$d^{k+1} \text{ points into } \sigma_{k+1} \text{ from } x^{k+1},$$
$$\text{and } \sigma_{k+1} \in \mathcal{M} \backslash \sigma_k.$$

9) goto next k.

FIGURE 1. Algorithm AGE

2. At step 8 in the algorithm, there may be more than one possible choice of cells σ_{k+1}. However, a lexicographic ordering, as described by Eaves (1976, §15), can be used to resolve any ambiguity concerning which cell to choose. The use of such a lexicographic ordering will be assumed in the convergence proof, and will be presented in more detail in the discussion preceding Lemma 2.9.

3. The requirement that $\|d^{k+1}\| = 1$ is arbitrarily chosen to force the choice of d^{k+1} to be unique.

4. The requirement that $x^1 - \mu d^1 \in \mathrm{int}(\sigma_1), \ \forall \mu \geq 0$ guarantees that the zero curve of $F(x, \mu) := G(x) + \mu B d^1$ contains a ray, and therefore assures us that it will not have any loops. This fact will be useful in our convergence proof. However, we shall also show that, under the assumption of coherent orientation, v_k is always negative, which by itself guarantees that no loops occur. Thus, under the assumption of coherent orientation, it is not necessary to find a ray start. However, in future work, we will prove convergence for a broader class of problems, in which case the ray start requirement will be useful.

The next few pages are devoted to proving the following convergence theorem:

THEOREM 2.5. *Let $\mathcal{M}$ be a finite collection of n-cells whose relative interiors are disjoint*

and whose union is $\mathbf{R}^n$. Let G be a coherently oriented, $\mathcal{M}$-PA function. Algorithm AGE, using lexicographic ordering, terminates after finitely many steps with a zero x^* of G.

PROOF (Outline).   There are three main parts to the proof. First, as Lemma 2.9, we will show that the algorithm terminates at a solution if $\mathcal{M}$ is a *subdivision* of $\mathbf{R}^n$ (see Definition 2.6). Second, in Lemma 2.11, we will show that even if $\mathcal{M}$ is not a subdivision of $\mathbf{R}^n$, there is a *refinement* (see Definition 2.10) $\mathcal{N}$ of $\mathcal{M}$ that is a subdivision. Finally, we show in Lemma 2.12 that if a subdivision $\mathcal{N}$ is a refinement of $\mathcal{M}$, then running the algorithm using $\mathcal{N}$ will generate exactly the same path as would be generated by using $\mathcal{M}$. Thus, the fact that the algorithm terminates at a solution using $\mathcal{N}$ guarantees that it will terminate at a solution using $\mathcal{M}$.   □

We now prove the three lemmas mentioned above. At this point, we recommend that the impatient reader skip ahead to Theorem 2.13. Our proof technique is based on the work of Eaves (1976). Eaves' analysis relies heavily on the notion of a subdivided manifold:

DEFINITION 2.6 (subdivided manifold).   Let N be a set in some Euclidean space, and let $\mathcal{N}$ be a finite or countable collection of n-cells in that space such that $N = \bigcup_{\sigma \in \mathcal{N}} \sigma$. Let $\tilde{\mathcal{N}}$ be the collection of all faces of elements of $\mathcal{N}$. $(N, \tilde{\mathcal{N}})$ is a subdivided n-manifold if

1. any two n-cells of $\mathcal{N}$ are either disjoint or meet in a common face;
2. each point of N has a neighborhood meeting only finitely many n-cells of $\mathcal{N}$;
3. each $(n - 1)$-cell of $\tilde{\mathcal{N}}$ lies in at most two n-cells.

If $(N, \tilde{\mathcal{N}})$ is a subdivided n-manifold for some subdivision $\mathcal{N}$, we call N an n-manifold and we call $\mathcal{N}$ a subdivision of N.

The following lemma shows that when $N = \mathbf{R}^n$, item 3 in Definition 2.6 is redundant. This result was proved by Robinson (1992) in the proof of Proposition 2.4. While Robinson's proposition is stated for the normal manifold, his proof is valid for general subdivisions of $\mathbf{R}^n$.

LEMMA 2.7.   *If $\mathcal{N}$ is a collection of cells whose union is $\mathbf{R}^n$ and if $\mathcal{N}$ satisfies 1 and 2 of Definition 2.6, then $\mathcal{N}$ is a subdivision of $\mathbf{R}^n$.*

The next step in our analysis is to prove that the algorithm works whenever $\mathcal{M}$ is a subdivision of $\mathbf{R}^n$. In this case, by defining $\mathcal{S} := \{\sigma \times \mathbf{R}_+ | \sigma \in \mathcal{M}\}$, we see that $\mathcal{S}$ is a subdivision of $\mathbf{R}^n \times \mathbf{R}_+$ and further that F is $\mathcal{S}$-PA. The starting point (x^1, μ_1) of the algorithm lies interior to the cell $\eta_1 := \sigma_1 \times \mathbf{R}_+$ of $\mathcal{S}$. Further, the ray $\{(x^1, \mu_1) - \mu(d^1, -1) | \mu \geq 0\}$ lies within η_1. Let $\tilde{\mathcal{S}}$ be the collection of all faces of elements of $\mathcal{S}$. Algorithm AGE is then seen to be equivalent to the algorithm described by Eaves (1976, §10.2), with the following relationships between the algorithms:

Eaves' Algorithm	Algorithm AGE
$\mathcal{M}$	$\tilde{\mathcal{S}}$
$F(x)$	$F(x, \mu) := G(x) + \mu B^1 d^1$
x_k	(x^k, μ_k)
σ_k	η_k
v_k	(d^k, v_k)

To discuss the behavior of this algorithm in more detail, we need some definitions from Eaves (1976).

DEFINITION 2.8 (regularity).   Let $(N, \tilde{\mathcal{N}})$ be a subdivided $(n + 1)$-manifold, let $\mathcal{N}$ be the collection of n-cells in $\tilde{\mathcal{N}}$, and let $F : N \to \mathbf{R}^n$ be a $\mathcal{N}$-PA map. A point x in N is said to be degenerate (otherwise regular) if x lies in a cell σ of $\tilde{\mathcal{N}}$ with $\dim(F(\sigma)) < n$. A value y in $F(N)$ is said to be a degenerate value (otherwise a regular value) if $F^{-1}(y)$ contains a degenerate point.

Note that if y is a regular value, then $F^{-1}(y)$ cannot intersect any k-cells of $\mathcal{S}$ with $k < n$.

By the assumption of coherent orientation, G is one-to-one in every n-cell of $\mathcal{M}$. Thus, $\dim(F(\eta)) = n$ for all $(n + 1)$-cells η of $\mathcal{S}$. Since the starting point (x^1, μ_1) of the algorithm is interior to η_1, it is a regular point of F. According to Eaves (1976, Theorem 15.13), since $\mathcal{S}$ is finite, the algorithm generates, in finitely many steps, either a point (x^*, μ_*) in the boundary of $\mathbf{R}^n \times \mathbf{R}_+$, or a ray in $F^{-1}(0)$ different from the starting ray. In the first case, we know that $\mu_* = 0$, since the boundary of $\mathbf{R}^n \times \mathbf{R}_+$ is $\mathbf{R}^n \times \{0\}$. It then follows, from our earlier remarks, that x^* satisfies $G(x^*) = 0$. Therefore, to guarantee that the algorithm finds a solution, we need only show that it cannot produce a ray different from the starting ray.

We first consider the case when 0 is a *regular value* of F. In this case, by Eaves (1976, Theorem 9.1), $F^{-1}(0)$ is a 1-manifold which is subdivided by sets of the form $\eta \cap F^{-1}(0)$. Further, since $F^{-1}(0)$ cannot intersect any k-cells with $k < n$, each point on $F^{-1}(0)$ is in at most two $(n + 1)$-cells of $\mathcal{S}$. Thus, in step 8 of the algorithm, the choice of σ_{k+1} is well-defined. (The only difficulty would be if (x^{k+1}, μ_{k+1}) lies in only one $(n + 1)$-cell μ_k so that no σ_{k+1} could be selected. But in this case, (x^{k+1}, μ_{k+1}) would be a boundary point of $\mathbf{R}^n \times \mathbf{R}_+$. Thus, $\mu_{k+1} = 0$, so the algorithm would have terminated in step 5.)

Let (d^k, v_k) be the direction of the path within the $(n + 1)$-cell η_k of $\mathcal{S}$, and let G have representation (B^k, b^k) on the n-cell σ_k of $\mathcal{M}$. Then by Eaves (1976, Lemma 12.3), the *curve index*, given by

$$(\operatorname{sgn} v_k)(\operatorname{sgn} \det B^k)$$

is constant everywhere along the path. Since $v_1 = -1$ for the starting direction (d^1, v_1), and since G is coherently oriented, it follows that v_k is negative in each cell that the path enters. But this means that the parameter μ decreases strictly in each cell. Thus, after finitely many steps, we must have $\mu = 0$.

When 0 is a degenerate value of F, $F^{-1}(0)$ may intersect a k-cell of $\mathcal{S}$ with $k < n$. Thus, in step 8 of the algorithm, there may be multiple choices for which cell σ_{k+1} to enter next. To address this problem, a lexicographic ordering can be used to resolve ambiguities concerning which cell the path will enter. Such a scheme is conceptually equivalent to solving a perturbed problem, which we now describe.

Let $X = [\xi^1, \ldots, \xi^n]$ be an $(n + 1) \times n$ matrix such that $[X, (d^1; -1)]$ is of rank $(n + 1)$. Define the vector $[\epsilon] := (\epsilon^1, \epsilon^2, \ldots, \epsilon^n)^\top$ (note: the superscripts here refer to exponentiation). Define $(x^1(\epsilon); \mu_1(\epsilon)) := (x^1; \mu_1) + X[\epsilon]$. Since $(x^1; \mu_1)$ in Algorithm AGE is interior to η_1, then $(x^1(\epsilon); \mu_1(\epsilon)) \in \operatorname{int}(\mu_1)$ for small enough ϵ. Further, since $(-d^1, 1) \in \operatorname{int}(\operatorname{rec}(\mu_1))$, $(x^1(\epsilon, \mu_1(\epsilon)) + \mu(-d^1, 1) \in \operatorname{int}(\eta_1)$, for all $\mu \geq 0$. Thus, $x^1(\epsilon)$, $\mu_1(\epsilon)$, σ_1, and d^1 satisfy the starting conditions needed to apply the algorithm to the perturbed problem given by

$$0 = F(x, \mu) - p(\epsilon),$$

where $p(\epsilon) := F(x^1(\epsilon), \mu_1(\epsilon))$. Observe that

$$p(\epsilon) = F(x^1, \mu_1) + [B^1, B^1 d^1]X[\epsilon]$$

(6)
$$= [B^1, B^1 d^1]X[\epsilon]$$

$$= Y[\epsilon],$$

where $Y := [B^1, B^1 d^1]X$. Y is an invertible $n \times n$ matrix, so that by Eaves (1976, Lemma 14.2), $p(\epsilon)$ is a regular value for all ϵ sufficiently small. Thus, by the arguments given above

for regular values, using Algorithm AGE to solve the perturbed problem will, after a finite number of steps J, produce a point $(x^J(\epsilon))$ such that $G(x^J(\epsilon)) = p(\epsilon)$.

Let $(x^k(\epsilon), \mu_k(\epsilon))$ be the sequence of points generated by the algorithm for the perturbed problem. By the discussion in Eaves (1976, §15), there is a sequence of matrices $X^k \in \mathbf{R}^{(n+1) \times n}$ and a sequence of points (x^k, μ_k) such that $(x^k(\epsilon); \mu_k(\epsilon)) := (x^k; \mu_k) + X^k[\epsilon]$ for all small ϵ. The points (x^k, μ_k) are exactly the sequence of points generated by the algorithm for solving the unperturbed problem using the lexicographic ordering. Since the algorithm terminates after J steps for all small ϵ, we see that $\mu_J(\epsilon) = 0$ and $G(x^J(\epsilon)) = p(\epsilon)$. It follows that $\mu_J = 0$ and further that $G(x^J) = 0$. Thus, using a lexicographic ordering, the algorithm finds a solution after a finite number of steps.

We have proven the following lemma:

LEMMA 2.9. *Let $\mathcal{M}$ be a subdivision of $\mathbf{R}^n$ and G be a coherently oriented, $\mathcal{M}$-PA function. Algorithm AGE, using lexicographic ordering, terminates after finitely many steps with a zero x^* of G.*

We now address the case where $\mathcal{M}$ is not a subdivision of $\mathbf{R}^n$. We begin by proving that $\mathcal{M}$ can be *refined* to produce a subdivision.

DEFINITION 2.10 (refinement). Let $\mathcal{M}$ and $\mathcal{N}$ be finite collections of n-cells. $\mathcal{N}$ is said to be a refinement of $\mathcal{M}$ if each cell σ of $\mathcal{M}$ is the union of a finite collection of cells τ_i of $\mathcal{N}$, and if each cell of $\mathcal{N}$ is contained in some cell of $\mathcal{M}$.

The following lemma is proved by Hudson (1969, Lemma 1.5), however, using different nomenclature. In particular, the term "subdivision" is used in place of our term "refinement."

LEMMA 2.11. *Let $\mathcal{M} := \{C_i\}$ be a collection of $J < \infty$ n-cells which covers $\mathbf{R}^n$. There exists a subdivision $\mathcal{N}$ of $\mathbf{R}^n$ such that $\mathcal{N}$ is a refinement of $\mathcal{M}$.*

We now show that using $\mathcal{N}$, the algorithm follows the same path as it would using $\mathcal{M}$.

LEMMA 2.12. *Let G be a coherently oriented $\mathcal{M}$-PA function, where $\mathcal{M}$ is a finite collection of relatively disjoint n-cells whose union is $\mathbf{R}^n$. Let $\mathcal{N}$ be a refinement of $\mathcal{M}$ such that $\mathcal{N}$ is a subdivision of $\mathbf{R}^n$. Then Algorithm AGE, using lexicographic ordering, will find a solution x^* to $G(x) = 0$ in a finite number of steps. Furthermore, the sequence of points generated by the algorithm using $\mathcal{M}$ is a subsequence of the points that would be generated using $\mathcal{N}$.*

PROOF. Consider first running the algorithm using $\mathcal{N}$ instead of $\mathcal{M}$. By Lemma 2.9, the algorithm will terminate after some finite number of steps J. The algorithm will visit a sequence of n-cells $\{\tau_k\} \subset \mathcal{N}$, $k = 1, \ldots, J$, and will generate a sequence of points $\{(x^k, \mu_k)\}$, $k = 1, \ldots, J + 1$ and directions $\{(d^k, v_k)\}$, $k = 1, \ldots, J$.

Let $j_1 := 1$ and let σ_1 be the unique cell in $\mathcal{M}$ that contains τ_1. Then for $i = 2, \ldots$, let j_i be the smallest index greater than j_{i-1} such that $\tau_{j_i} \not\subset \sigma_{i-1}$, setting $j_i = J + 1$ if no such index exists. If $j_i \leq J$, let σ_i be the unique cell in $\mathcal{M}$ that contains τ_{j_i}. Let K be such that $j_{K+1} = J + 1$. This process defines a sequence of cells $\{\sigma_i\}$, $i = 1, \ldots, K$ and indices $\{j_i\}$, $i = 1, \ldots, K + 1$, with the property that $\tau_k \subset \sigma_i$ whenever $j_i \leq k < j_{i+1}$, $i = 1, \ldots, K$.

We will show that if the algorithm is run using $\mathcal{M}$, then the sequence of points $\{(\xi^i, \nu_i)\}$ generated by the algorithm satisfies the equation $(\xi^i, \nu_i) = (x^{j_i}, \mu_{j_i})$, for $i = 1, \ldots, K + 1$. Thus, $(\xi^{K+1}, \nu_{K+1}) = (x^{J+1}, \mu_{J+1})$, so the algorithm finds a solution after a finite number of steps.

Let $\{(\delta^i, \zeta_i)\}$ be the sequence of directions chosen by the algorithm using $\mathcal{M}$. Clearly, since the algorithm is started at the point (x^1, μ_1) in the direction $(d^1, -1)$, the following is true: $(\xi^1, \nu_1) = (x^{j_1}, \nu_{j_1})$, $(\delta^1, \zeta_1) = (d^{j_1}, v_{j_1})$, and the first cell visited by the algorithm is σ_1.

We now proceed by induction: Assume that $(\xi^i, \nu_i) = (x^{j_i}, \nu_{j_i})$, $(\delta^i, \zeta_i) = (d^{j_i}, v_{j_i})$, and

that, using $\mathcal{M}$, the ith cell visited by the algorithm is σ_i. We shall prove that $(\xi^{i+1}, \nu_{i+1}) = (x^{j_{i+1}}, \nu_{j_{i+1}})$, $(\delta^{i+1}, \zeta_{i+1}) = (d^{j_{i+1}}, \nu_{j_{i+1}})$, and that the $(i + 1)$st cell visited by the algorithm is σ_{i+1}.

Let (B^i, b^i) be the representation of G on σ_i. This is also the representation of G on τ_k whenever $j_i \leq k < j_{i+1}$. Thus, in step 8 of the algorithm using $\mathcal{N}$, the direction (d^{k+1}, ν_{k+1}) chosen when entering cell τ_{k+1} must satisfy

$$B^i d^{k+1} + \nu_{k+1} B^1 d^1 = 0$$

for $j_i \leq k < j_{i+1}$. Since G is coherently oriented, B^i is invertible. Further, $\|d^{k+1}\| = 1$ and, by our earlier discussion, ν_{k+1} is negative. Thus, the direction is uniquely determined by the representation. In particular, $(d^{j_i}, \nu_{j_i}) = (d^{j_i+1}, \nu_{j_i+1}) = \cdots = (d^{j_{i+1}}, \nu_{j_{i+1}}) = (\delta^i, \zeta_i)$. From this it is clear that $x^{j_{i+1}}$ lies on the ray $\{\xi^i(\theta)|\theta \geq 0\}$, where $\xi^i(\theta) := \xi^i + \theta\delta^i$. Further, $x^{j_{i+1}}$ is on the boundary of σ_i.

If the ray $\{\xi^i(\theta)|\theta \geq 0\}$ contains a point in the interior of σ_i, then the ray cannot be extended past $x^{j_{i+1}}$ without exiting σ_i. Thus, $x^{j_{i+1}} = \xi^i(\theta_i)$ where $\theta_i := \sup\{\theta|\xi^i(\theta) \in \sigma_i\}$. In other words, $(x^{j_{i+1}}, \mu_{j_{i+1}}) = (\xi^{i+1}, \nu_{i+1})$.

If the ray $\{\xi^i(\theta)|\theta \geq 0\}$ does not contain an interior point of σ_i, then we must resort to the lexicographic ordering to prove that $x^{j_{i+1}} = \xi^i(\theta_i)$. Since σ_i and $\tau_{j_{i+1}}$ are relatively disjoint convex sets, there exists a separating hyperplane H_i defined by a vector c^i, and a scalar α_i such that $c^{i^\top}x < \alpha_i$, $\forall x \in \text{int}(\sigma_i)$, and $c^{i^\top}x \geq \alpha_i$, $\forall x \in \tau_{j_{i+1}}$. Suppose we run the algorithm using $\mathcal{N}$ to solve the perturbed problem $G_\epsilon(x) := G(x) - p(\epsilon) = 0$, where $p(\epsilon)$ is defined by (6). Then, for ϵ small enough, the algorithm will visit the same sequence of cells $\{\tau_k\}$ as it visits in the unperturbed problem. Also, by our earlier discussion, the algorithm will generate the sequence of points $\{(x^k(\epsilon); \mu_k(\epsilon))\} = \{(x^k; \mu_k) + X^k[\epsilon]\}$, where $\{X^k\}$ is a fixed sequence of matrices.

Since 0 is a regular value of G_ϵ, $\dim(G_\epsilon(\tau_k \cap \tau_{k+1})) \geq n - 1$ for any k. Thus, $G_\epsilon^{-1}(0)$ contains only one point in $\tau_k \cap \tau_{k+1}$, namely x^{k+1}. Therefore, the direction d^{k+1} must point into the interior of τ_{k+1}.

By similar arguments as before, $x^{j_{i+1}}$ lies on the ray $\{\xi_\epsilon^i(\theta)|\theta \geq 0\}$, where $\xi_\epsilon^i(\theta) := x^{j_i}(\epsilon) + \theta(d^{j_i})$. But, since d^{k+1} points into the interior of τ_{j_i}, this ray must contain a point $\hat{x}$ in the interior of σ_i. Thus, $c^{i^\top}\hat{x} < \alpha_i$. But $c^{i^\top}x^{j_{i+1}} \geq \alpha_i$ since $x^{j_{i+1}}$ is in $\tau_{j_{i+1}}$. It follows that $c^{i^\top}d^{j_i} > 0$. Thus, even for the unperturbed problem, the ray $\xi^i(\theta)$ cannot be extended past the point $x^{j_{i+1}}$ without crossing the hyperplane H_i, and thereby exiting σ_i. Thus, $x^{j_{i+1}} = \xi^i(\theta_i)$, and as before, $(x^{j_{i+1}}, \mu_{j_{i+1}}) = (\xi^{i+1}, \nu_{i+1})$.

Finally, note that for all small ϵ, the point $x^{j_{i+1}}(\epsilon)$ is a regular point, so $\tau_{(j_{i+1}-1)}$ and $\tau_{j_{i+1}}$ are the only n-cells of $\mathcal{N}$ that contain $x^{j_{i+1}}(\epsilon)$. Thus, σ_i and σ_{i+1} are the only n-cells of $\mathcal{M}$ that contain $x^{j_{i+1}}(\epsilon)$. Thus, for all small ϵ, the algorithm, using $\mathcal{M}$ will enter cell σ_{i+1} at the next iteration. But this means that using lexicographic ordering the algorithm will enter cell σ_{i+1} next when solving the unperturbed problem. Finally, since the representation of G on σ_{i+1} is identical to the representation of G on $\tau_{j_{i+1}}$, we must have $(\delta^{i+1}, \zeta_{i+1}) = (d^{j_{i+1}}, \nu^{j_{i+1}})$.

The lemma is now proved by induction. $\square$

This completes the proof of Theorem 2.5. Our final task in this section is to establish the claim made in Comment 1 following Algorithm AGE.

THEOREM 2.13. *Under the hypothesis of Theorem 2.5, let $\{\sigma_k\}$ be the sequence of cells chosen in Step 8 of Algorithm AGE using lexicographic ordering, and let (B^k, b^k) represent A_T on σ_k. Then $B^{k+1} - B^k$ has rank 1.*

PROOF. Using lexicographical ordering, the algorithm will choose the same cell σ_{k+1} in step 8 as it would when solving the perturbed problem for small ϵ. However, 0 is a regular value for the perturbed problem, so $\tau_k := \sigma_k \cap \sigma_{k+1}$ must have dimension $n - 1$. Now, for any two points $x^1, x^2 \in \tau_k$

$$B^{k+1}(x^1 - x^2) = B^k(x^1 - x^2) \Rightarrow (x^1 - x^2) \in \ker(B^{k+1} - B^k)$$

Thus, $\dim \ker(B^{k+1} - B^k) = n - 1$ so $\operatorname{rank}(B^{k+1} - B^k) = 1$. $\square$

3. The T-map. The T-map, denoted F_T, is a generalization of the normal map that is formed by replacing the projection operator π_C in (3) by the resolvent operator $P_T := (I + T)^{-1}$. Specifically, the T-map is given by

$$(7) \qquad F_T(x) := F(P_T(x)) + x - P_T(x).$$

We assume throughout that T is a maximal monotone operator. In this case, Minty (1962) showed that P_T is a continuous, single-valued, nonexpansive function defined on all of $\mathbf{R}^n$. Since the image of P_T is $\operatorname{dom}(T)$ (which is contained in the domain of F, Ω), it follows that F_T is a single-valued function defined on all of $\mathbf{R}^n$.

By Brézis (1973, Example 2.1.2), $I + T$ is monotone, and therefore P_T is monotone. We now show that solving GE(F, T) is equivalent to finding a zero of F_T.

THEOREM 3.1. *Given a maximal monotone multifunction $T : \mathbf{R}^n \rightrightarrows \mathbf{R}^n$, and a function $F : \Omega \subset \mathbf{R}^n \to \mathbf{R}^n$, let F_T be defined by (7). If x is a zero of F_T, then $z := P_T(x)$ solves GE(F, T). Conversely, if z solves GE(F, T), then $x := z - F(z)$ is a zero of F_T.*

PROOF. Suppose $F_T(x) = 0$ and let $z := P_T(x)$. Then $0 = F_T(x) = F(z) + x - z$, and

$$-F(z) = x - z$$

$$\in (I + T)(I + T)^{-1}(x) - z$$

$$= (I + T)(z) - z$$

$$= T(z).$$

Conversely, suppose $-F(z) \in T(z)$ and let $x := z - F(z)$. Then $x \in z + T(z) = (I + T)(z)$, so $P_T(x) = (I + T)^{-1}(x) = z$. Thus $F_T(x) = F(z) + x - z = F(z) - F(z) = 0$. $\square$

So far, we have not made any assumptions on T other than that it is maximal monotone. We now focus on the case where T is polyhedral.

DEFINITION 3.2 (polyhedral). A multifunction T is polyhedral if its graph is the union of finitely many polyhedral convex sets.

Our first task will be to show that, for polyhedral T, the resolvent operator $P_T := (I + T)^{-1}$ is a piecewise-affine map.

LEMMA 3.3. *A single-valued multifunction $T : \mathbf{R}^n \rightrightarrows \mathbf{R}^m$ whose graph is a convex polyhedron is affine on $\operatorname{dom}(T)$.*

PROOF. Assume $\operatorname{dom}(T) \neq \varnothing$. (Otherwise the lemma is true vacuously.) Since the graph of T is a polyhedron, T can be written as $T = \{(x, y) \mid Ax + By \geq c\}$ for some $A \in \mathbf{R}^{p \times n}$, $B \in \mathbf{R}^{p \times m}$, and $c \in \mathbf{R}^p$, where p is some nonnegative integer. Let $\mathcal{H} := \{i \mid A_{i\cdot}x + B_{i\cdot}y = c_i, \forall (x, y) \in T\}$. In words, $\mathcal{H}$ is the set of row indices for which the corresponding constraint is active for *all* points in T.

We first establish that $\mathcal{H}$ is nonempty and that $\ker B_{\mathcal{H}\cdot} = \{0\}$. To do this, let $\bar{\mathcal{H}} := \{i \mid i \notin \mathcal{H}\}$. If $\bar{\mathcal{H}}$ is nonempty, then for each $i \in \bar{\mathcal{H}}$, $\exists (x^i, y^i) \in T$ with $A_{i\cdot}x^i + B_{i\cdot}y^i > c_i$. Let $(\tilde{x}, \tilde{y}) = \sum_{i \in \bar{\mathcal{H}}} (x^i, y^i)/|\bar{\mathcal{H}}|$, where $|\bar{\mathcal{H}}|$ is the cardinality of the index set $\bar{\mathcal{H}}$. If $\bar{\mathcal{H}} = \varnothing$, choose

$(\tilde{x}, \tilde{y})$ arbitrarily in T. In either case, $(\tilde{x}, \tilde{y}) \in T$ and also, $A_{i.}\tilde{x} + B_{i.}\tilde{y} > c_i$, $\forall i \in \mathcal{H}$. Let $\hat{y}$ be any nonzero vector such that $B_{i.}\hat{y} = 0$ for all $i \in \mathcal{H}$. Then, for $\epsilon > 0$ small enough, $A\tilde{x} + B(\tilde{y} + \epsilon\hat{y}) \geq c$. Thus, $(\tilde{x}, \tilde{y} + \epsilon\hat{y}) \in T$, contradicting the single-valuedness of T. It follows that $\mathcal{H} \neq \varnothing$ and that ker $B_{\mathcal{H}.} = \{0\}$.

Now, by the definition of $\mathcal{H}$, we have $(x, y) \in T \Rightarrow A_{\mathcal{H}.}x + B_{\mathcal{H}.}y = c_{\mathcal{H}}$. Conversely, suppose (x, y) satisfies $A_{\mathcal{H}.}x + B_{\mathcal{H}.}y = c_{\mathcal{H}}$. If $x \in \text{dom}(T)$, then $\exists \hat{y}$ such that $(x, \hat{y}) \in T$. But this means that $A_{\mathcal{H}.}x + B_{\mathcal{H}.}\hat{y} = c_{\mathcal{H}}$, which implies that $\hat{y} - y \in \text{ker } B_{\mathcal{H}.} = \{0\}$. That is $\hat{y} = y$. We have thus shown that

$$(x, y) \in T \Leftrightarrow x \in \text{dom}(T) \quad \text{and} \quad A_{\mathcal{H}.}x + B_{\mathcal{H}.}y = c_{\mathcal{H}}.$$

Finally, since ker $B_{\mathcal{H}.} = \{0\}$, $B_{\mathcal{H}.}$ has a left inverse $R \in \mathbf{R}^{m \times p}$. Thus, for $x \in \text{dom}(T)$,

$$(x, y) \in T \Leftrightarrow A_{\mathcal{H}.}x + B_{\mathcal{H}.}y = c_{\mathcal{H}}$$

$$\Leftrightarrow y = Rc_{\mathcal{H}} - RA_{\mathcal{H}.}x.$$

So T is an affine function on $\text{dom}(T)$. $\quad\square$

THEOREM 3.4. *Given a maximal monotone polyhedral multifunction $T : \mathbf{R}^n \rightrightarrows \mathbf{R}^n$, the resolvent operator $P_T := (I + T)^{-1}$ is a piecewise affine function on all of $\mathbf{R}^n$.*

PROOF. Since T is polyhedral, $I + T$ is also polyhedral (Robinson 1979a) and therefore so is $P_T = (I + T)^{-1}$. Thus, $P_T = \cup \Gamma_i$, where $\{\Gamma_i\}$ is a finite collection of polyhedral convex sets. Let C_i be the projection of Γ_i onto the domain of P_T (i.e., $C_i = \pi_1(\Gamma_i)$, where $\pi_1 := (x, y) \mapsto x$). Define $\mathcal{M} := \{C_i | \dim(C_i) = n\}$.

Since P_T is defined on all of $\mathbf{R}^n$, $\cup C_i = \mathbf{R}^n$. Let $M := \cup_{C_i \in \mathcal{M}} C_i$. Since M is closed, its complement, $\backslash M := \mathbf{R}^n \backslash M$, is open. Thus, $\backslash M$ is either the empty set, or it has nonempty interior. But $\backslash M \subset \cup_{\dim(C_i) < n} C_i$. Thus, $\backslash M$ has no interior. In other words $\backslash M = \varnothing$ and thus, $M := \mathbf{R}^n$.

To show that P_T is $\mathcal{M}$-PA, all that is needed is to show that for each $C_i \in \mathcal{M}$, the restriction of P_T to C_i is affine. However, since P_T is single-valued, the graph of P_T restricted to C_i is simply the convex polyhedral set Γ_i. By Lemma 3.3, P_T is affine on C_i. $\quad\square$

COROLLARY 3.5. *If T is a maximal monotone polyhedral multifunction and F is affine, then the T-map, F_T, defined by (7) is piecewise affine.*

4. Affine generalized equations. We now show how to apply the algorithm of §2 to construct an algorithm to solve the *affine generalized equation*:

$$(8) \qquad\qquad 0 \in Ax - a + T(x),$$

where $A \in \mathbf{R}^{n \times n}$, $a \in \mathbf{R}^n$, and T is a maximal monotone polyhedral multifunction. For this problem, the T-map is given by

$$(9) \qquad\qquad A_T := AP_T(x) + x - P_T(x) - a.$$

As was shown in §3, for polyhedral T, A_T is piecewise affine with respect to some *finite* collection $\mathcal{M}$ of n-cells whose union is $\mathbf{R}^n$. Thus, to complete the description of the algorithm for affine generalized equations, it remains to show how to generate the representations.

The task of constructing $\mathcal{M}$ is dependent upon how T is described. For example, in Robinson (1992), T is taken as the normal cone N_C to a polyhedral convex set C. $\mathcal{M}$ is then

chosen to be the *normal manifold*, which is defined in terms of the nonempty faces F_i of C. Specifically, the cells of the normal manifold are defined by

$$\sigma_i := F_i + N_{F_i},$$

where N_{F_i} is the common value of N_{F_i} for $x \in \mathrm{ri}(F_i)$. This particular choice of cells leads to the algorithm given in Cao and Ferris (1996b).

For more general T, we assume that T is described as the union of a finite collection of polyhedral convex sets C_i. We can then describe P_T as the union of the sets $S_i := \{(x + y, x) | (x, y) \in C_i\}$. By projecting each S_i onto the domain of P_T, we produce a collection of sets

$$\sigma_i := \{x + y | (x, y) \in C_i\}.$$

Since we know $\mathrm{dom}(P_T) = \mathbf{R}^n$, it follows that $\cup \sigma_i = \mathbf{R}^n$ and thus $\mathcal{M} := \{\sigma_i | \mathrm{int}(\sigma_i) \neq \varnothing\}$.

To provide an example of this process, we return to the case where $T = N_C$. Observe that

$$N_C = \bigcup_{F_i} \bigcup_{x \in F_i} \{x\} \times N_{F_i} = \bigcup_{F_i} F_i \times N_{F_i}.$$

Thus, we see that $\mathbf{R}^n$ is the union of the polyhedral convex sets

$$C_i := \{x + y | (x, y) \in F_i \times N_{F_i}\} = F_i + N_{F_i}.$$

It follows that the process described above yields the normal manifold.

Robinson (1992, Proposition 2.4) proved that the normal manifold is a subdivision of $\mathbf{R}^n$. However, in general, the collection of cells $\mathcal{M}$ generated by the above process is not a subdivision. This can be demonstrated by the following example. Let

$$C_1 := \{(x, 0) \in \mathbf{R}^2 \times \mathbf{R}^2 | x_1 \leq 0\},$$

$$C_2 := \{(x, 0) \in \mathbf{R}^2 \times \mathbf{R}^2 | x_1 \geq 0, x_2 \geq 0\},$$

$$C_3 := \{(x, 0) \in \mathbf{R}^2 \times \mathbf{R}^2 | x_1 \geq 0, x_2 \leq 0\},$$

and let $T := \cup_{i=1}^3 C_i$. Observe that T is simply the zero mapping, and is thus a maximal monotone multifunction. However, employing our procedure for constructing $\mathcal{M}$, we obtain $\sigma_1 = \{x \in \mathbf{R}^2 | x_1 \leq 0\}$, $\sigma_2 = \{x \in \mathbf{R}^2 | x_1 \geq 0, x_2 \geq 0\}$, $\sigma_3 = \{x \in \mathbf{R}^2 | x_1 \geq 0, x_2 \leq 0\}$. Since $\sigma_1 \cap \sigma_2$ is not a face of σ_1, we see that $\mathcal{M} := \{\sigma_1, \sigma_2, \sigma_3\}$ is not a subdivision of $\mathbf{R}^n$.

Since P_T is single-valued, then by Lemma 3.3, P_T is affine on each cell $\sigma_i \in \mathcal{M}$. A representation of A_T on each cell is then given by (9). In order to have a workable description of these affine maps, it would appear necessary to exploit the underlying structure of T. One such case is the subject of the next section.

5. Separable T. A particularly important class of affine variational inequalities is that for which the set C is *rectangular*, i.e., C is defined by the constraints

$$l \leq z \leq u,$$

where l and u are vectors in $\mathbf{R}^n$, with $l_i \in [-\infty, \infty)$ and $u_i \in (-\infty, \infty]$ for $1 \le i \le n$. This problem class has a number of features that are very attractive for pivotal algorithms similar to Algorithm AGE. In particular, the cells of linearity of the normal map are rectangular, and furthermore the normal map itself takes on a very simple form. Specifically, for an affine function $F(z) := Az + b$, the matrix used to represent the normal map on any cell is formed simply by replacing some of the columns of A by the corresponding columns of the identity matrix.

Rectangular variational inequalities are also attractive from a theoretical standpoint. In particular, if at least one of l_i and u_i is finite for each i, then the normal map is coherently oriented with respect to C if and only if A is a *P-matrix*.

DEFINITION 5.1 ((Cottle, Pang and Stone 1992)). A matrix A is said to be a *P*-matrix if all its principal minors are positive.

Note that when C is rectangular, then $N_C(z) = N_{[l,u]}(z) = \prod_{i=1}^n N_{[l_i, u_i]}(z_i)$. This suggests that we can extend the notion of rectangularity to generalized equations by requiring that the multifunction T be *separable*, i.e., it is of the form

$$T(z) = \begin{bmatrix} T_1(z_1) \\ T_2(z_2) \\ \vdots \\ T_n(z_n) \end{bmatrix},$$

where for each i, T_i is a maximal monotone polyhedral multifunction from $\mathbf{R}$ to $\mathbf{R}$. With such a T, we shall see that the cells of linearity of the T-map A_T are rectangular.

We begin by looking at the resolvent operator $P_T = (I + T)^{-1}$. Note that

$$P_T(x) = \begin{bmatrix} P_{T_1}(x_1) \\ P_{T_2}(x_2) \\ \vdots \\ P_{T_n}(x_n) \end{bmatrix},$$

where for each i, $P_{T_i} = (1 + T_i)^{-1}$. Since P_T is a continuous piecewise affine function, it follows that P_{T_i} is continuous piecewise affine function from $\mathbf{R}$ into $\mathbf{R}$. Let k_i be the number of breakpoints of P_{T_i}. Then, for some strictly increasing sequence of breakpoints $\{\xi_{ij}\}$, $j = 1, \ldots, k_i$ and some set of coefficients $\{d_{ij}, b_{ij}\}$, $j = 0, \ldots, k_i$,

$$P_{T_i}(x) = \begin{cases} d_{i0}x + b_{i0} & x \le \xi_{i1}, \\ d_{ij}x + b_{ij} & \xi_{ij} \le x \le \xi_{i(j+1)}, \ 1 \le j < k_i, \\ d_{ik_i}x + b_{ik_i} & \xi_{ik_i} \le x. \end{cases}$$

Note that since P_T is monotonic and nonexpansive, $0 \le d_{ij} \le 1$.

The breakpoint sequence defines a subdivision of $\mathbf{R}$ given by $\mathcal{M}_i = \bigcup_{j=0}^{k_i} \sigma_{ij}$, where

$$\sigma_{ij} = \left\{ x \ \middle| \ \begin{array}{ll} x \le \xi_{i1} & j = 0 \\ \xi_{ij} \le x \le \xi_{i(j+1)} & 0 < j < k_i \\ \xi_{ik_i} \le x & j = k_i \end{array} \right\}.$$

We then define a subdivision of $\mathbf{R}^n$ by $\mathcal{M} = \prod_{i=1}^n \mathcal{M}_i$ with n-cells defined by $\sigma_{[j1, j2, \ldots, jn]} = \sigma_{1j_1} \times \sigma_{2j_2} \times \cdots \times \sigma_{nj_n}$.

Clearly, P_T (and therefore A_T) is $\mathcal{M}$-PA. This establishes our earlier statement that the cells of linearity of A_T are rectangular. For each cell $\sigma_{[j_1, j_2, \dots, j_n]}$ of $\mathcal{M}$, define a diagonal matrix $D_{[j_1, j_2, \dots, j_n]}$ by $D_{[j_1, j_2, \dots, j_n]}(i, i) = d_{ij_i}$. Further, define the vector $b_{[j_1, j_2, \dots, j_n]} = [b_{1j_1}; b_{2j_2}; \dots ; b_{nj_n}]$. Then on $\sigma_{[j_1, j_2, \dots, j_n]}$, P_T is represented by $(D_{[j_1, j_2, \dots, j_n]}, b_{[j_1, j_2, \dots, j_n]})$. Thus, on $\sigma_{[j_1, j_2, \dots, j_n]}$, the T-map is given by

$$A_T(x) = A(P_T(x)) + a + x - P_T(x)$$

$$= (AD_{[j_1, j_2, \dots, j_n]} + I - D_{[j_1, j_2, \dots, j_n]})x + (A - I)b_{[j_1, j_2, \dots, j_n]} + a$$

$$= [d_{1j_1}A_{.1} + (1 - d_{1j_1})I_{.1}, d_{2j_2}A_{.2} + (1 - d_{2j_2})I_{.2}, \dots, d_{nj_n}A_{.n} + (1 - d_{nj_n})I_{.n}]x$$

$$\quad + (A - I)b_{[j_1, j_2, \dots, j_n]} + a$$

$$=: M_{[j_1, j_2, \dots, j_n]}x + b_{[j_1, j_2, \dots, j_n]}.$$

Thus, we see that the matrix $M_{[j_1, j_2, \dots, j_n]}$ which represents A_T on $\sigma_{[j_1, j_2, \dots, j_n]}$ has columns which are convex combinations of columns of A and the corresponding columns of I.

We now set about proving the main result of this section. Namely, if A is a P-matrix, then A_T is coherently oriented for any separable polyhedral maximal monotone multifunction T. We first need to prove two technical lemmas.

LEMMA 5.2.   *If A and B are $n \times n$ matrices where B is rank-1 such that $\det(A) > 0$ and $\det(A + B) > 0$, then $\det(A + \lambda B) > 0$ for all $\lambda \in [0, 1]$.*

PROOF.

$$\det(A + \lambda B) = \sum_{(j_1, \dots, j_n)} \det[C_{j_1}, \dots, C_{j_n}],$$

where the summation is taken over all possible choices of $(j_1, \dots, j_n)$ such that C_{j_i} is either $A_{.i}$ or $\lambda B_{.i}$. Since B is rank-1, the determinants in the above sum are zero for all choices that include at least two columns of λB. Thus,

$$\det(A + \lambda B) = \det A + \sum_{i=1}^{n} \det[A_{.1}, \dots, \lambda B_{.i}, \dots, A_{.n}]$$

$$= \det A + \lambda \left(\sum_{i=1}^{n} \det[A_{.1}, \dots, B_{.i}, \dots, A_{.n}] \right).$$

Thus, $\det(A + \lambda B)$ is an affine function of λ, which is positive at $\lambda = 0$ and $\lambda = 1$. Thus, it is positive for all $\lambda \in [0, 1]$. $\square$

LEMMA 5.3.   *Let A be an $n \times n$ matrix and let $\{B^1, \dots, B^k\}$ be a collection of rank-1 $n \times n$ matrices. If $\det(A + \lambda_1 B^1 + \cdots + \lambda_k B^k) > 0$ for all choices of $\lambda_i = 0$ or 1, then $\det(A + \lambda_1 B^1 + \cdots + \lambda_k B^k) > 0$ for all choices of $\lambda_i \in [0, 1]$.*

PROOF (By induction).   The lemma is true for $k = 1$ by Lemma 5.2. Now, suppose the lemma is true for all $k < m$, we shall prove the lemma true for $k = m$.

Suppose $\{B^1, \dots, B^m\}$ is a collection of rank-1 $n \times n$ matrices such that $\det(A + \lambda_1 B^1$

$+ \cdots + \lambda_m B^m) > 0$ for all choices of $\lambda_i = 0$ or 1. Let $\hat{A} := A + B^m$. Then $\hat{A}$ and $\{B^1, \ldots, B^{m-1}\}$ satisfy the conditions of the lemma for $k = m - 1$. Thus, if $\lambda_i \in [0, 1]$, $\forall i$, then

$$\det((A + B^m) + \lambda_1 B^1 + \cdots + \lambda_{m-1} B^{m-1}) > 0.$$

Similarly, with $\hat{A} := A$, we have

$$\det(A + \lambda_1 B^1 + \cdots + \lambda_{m-1} B^{m-1}) > 0.$$

From these two results, we see that if we let $\tilde{A} := A + \lambda_1 B^1 + \cdots + \lambda_{m-1} B^{m-1}$, then $\tilde{A}$ and B^m satisfy the hypotheses for Lemma 5.2. Thus, for $\lambda_m \in [0, 1]$,

$$0 < \det(\tilde{A} + \lambda_m B^m),$$

$$= \det(A + \lambda_1 B^1 + \cdots + \lambda_m B^m). \quad \square$$

THEOREM 5.4. *If A is a P-matrix, then for* any *separable maximal monotone polyhedral multifunction T, the T-map A_T defined by* (9) *has the property that in any cell of linearity, the matrix representing A_T has positive determinant. In particular, A_T is coherently oriented.*

PROOF. Let A_T have the representation $(\hat{A}, b)$ in the n-cell σ. By the our earlier discussion, $\hat{A}$ can be formed by replacing columns of A by a convex combination of columns of A and the corresponding columns of the identity matrix. Thus, the matrix is of the form

$$\hat{A} = A + \lambda_1 B^1 + \cdots + \lambda_n B^n, \qquad \lambda_i \in [0, 1]$$

where $B^i := (I_{\cdot i} - A_{\cdot i}) I_{\cdot i}^\top$. Observe that B^i is a rank-1 matrix.

Since A is a P-matrix, the matrix formed by replacing an arbitrary set of columns of A by corresponding columns of the identity matrix has positive determinant. Thus, the matrices A, $B^1, \ldots, B^n$ satisfy the hypotheses of Lemma 5.3. Thus, by Lemma 5.3, $\det(\hat{A}) > 0$. $\quad \square$

COROLLARY 5.5. *If A is a P-matrix and T is a separable maximal monotone polyhedral multifunction, then using lexicographic ordering, Algorithm AGE will find a solution to $A_T(x) = 0$ in a finite number of steps.*

6. Piecewise linear-quadratic programming.

We conclude by giving an example of a well known problem in mathematical programming that can be solved using the technique we have presented. The piecewise linear-quadratic programming problem (PLQP) is given by

$$(10) \qquad \min h(x) = f(x) + \phi(Ax),$$

where $A \in \mathbf{R}^{m \times n}$, and $f : \mathbf{R}^n \to \mathbf{R} \cup \{\infty\}$ and $\phi : \mathbf{R}^m \to \mathbf{R} \cup \{\infty\}$ are convex *piecewise linear-quadratic* functions, defined below.

DEFINITION 6.1. A function $f : \mathbf{R}^n \to \mathbf{R} \cup \{\infty\}$ is piecewise linear-quadratic if dom f is closed and convex and there exists a finite subdivision $\mathcal{M}$ of dom(f) such that for each $\sigma \in \mathcal{M}$, $f_{|\sigma}$ is a quadratic function.

Note that dom f is polyhedral, and further that since the cells in the subdivision are closed, f is a continuous function on dom f.

The optimality conditions for PLQP are stated by the relation

$$0 \in \partial h(x),$$

where ∂h is the *convex subdifferential* operator defined by

$$\partial h(x) := \{z | h(w) \geq h(x) + z^\top (w - x), \ \forall \, w \in \text{dom}(h)\}.$$

Under an appropriate constraint qualification (i.e. $\text{ri}(A(\text{dom}(f))) \cap \text{ri}(\text{dom}(\phi)) \neq \varnothing$), it follows that

$$\partial h(x) = \partial f(x) + A^\top \partial \phi(Ax).$$

Thus, for the optimality conditions to be satisfied, there must be an $x \in \text{dom}(f)$ and $y \in \partial \phi(Ax)$ such that $-A^\top y \in \partial f(x)$. By Rockafellar (1970, Theorem 23.5), the first statement is equivalent to

$$Ax \in \partial \phi^*(y),$$

where ϕ^* is the conjugate of ϕ. The optimality conditions are then

$$-A^\top y \in \partial f(x),$$

$$Ax \in \partial \phi^*(y).$$

Thus, if we define

$$\hat{A} := \begin{bmatrix} 0 & A^\top \\ -A & 0 \end{bmatrix}, \qquad T(x; y) := \begin{bmatrix} \partial f(x) \\ \partial \phi^*(y) \end{bmatrix},$$

the optimality conditions for PLQP can stated as the generalized equation

$$(11) \qquad\qquad -\hat{A}(x; y) \in T(x; y).$$

The fact that T is polyhedral was shown in Sun (1986). Thus, the optimality conditions for the piecewise linear-quadratic program can be expressed as an affine generalized equation, which can then be solved using our algorithm.

References

Aubin, J.-P., I. Ekeland. 1984. *Applied nonlinear analysis.* John Wiley & Sons, New York.

Brézis, H. 1973. *Opérateurs Maximaux Monotones et Semi-Groupes de Contractions dans les Espaces de Hilbert.* North-Holland, Amsterdam.

Cao, M., M. C. Ferris. 1996a. P_C matrices and the linear complementarity problem. *Linear Algebra and Its Applications* **246** 299–312.

———, ———. 1996b. A pivotal method for affine variational inequalities. *Math. Oper. Res.* **21** 44–64.

Cottle, R. W., J. S. Pang, R. E. Stone. 1992. *The Linear Complementarity Problem.* Academic Press, Boston, MA.

Dirkse, S. P. 1994. Robust solution of mixed complementarity problems. PhD thesis, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin. Available from ftp://ftp.cs.wisc.edu/math-prog/tech-reports/.

——, M. C. Ferris. 1995. The PATH solver: A non-monotone stabilization scheme for mixed complementarity problems. *Optimization Methods and Software* **5** 123–156.

——, ——. 1996. A pathsearch damped Newton method for computing general equilibria. *Ann. Oper. Res.* **68** 211–232.

Eaves, B. C. 1971. On the basic theorem of complementarity. *Math. Programming* **1** 68–87.

——. 1976. A short course in solving equations with PL homotopies. R. W. Cottle and C. E. Lemke, eds. *Nonlinear Programming.* American Mathematical Society, SIAM-AMS Proceedings, Providence, RI. 73–143.

——. 1978a. Computing stationary points. *Math. Programming Study* **7** 1–14.

——. 1978b. Computing stationary points, again. O. L. Mangasarian, R. R. Meyer and S. M. Robinson, eds. *Nonlinear Programming*, Vol. 3. Academic Press, London. 391–405.

——, C. E. Lemke. 1981. Equivalence of LCP and PLS. *Math. Oper. Res.* **6** 475–484.

Ferris, M. C., J. S. Pang. 1997. Engineering and economic applications of complementarity problems. *SIAM Rev.* **39** 669–713.

Hudson, J. F. P. 1969. *Piecewise Linear Topology.* Benjamin, New York.

Karamardian, S. 1971. Generalized complementarity problem. *J. Optimization Theory and Applications* **8** 161–167.

Lemke, C. E., J. T. Howson. 1964. Equilibrium points of bimatrix games. *SIAM J. Applied Math.* **12** 413–423.

Minty, G. J. 1962. Monotone (nonlinear) operators in Hilbert space. *Duke Math. J.* **29** 341–346.

Ralph, D. 1994. Global convergence of damped Newton's method for nonsmooth equations, via the path search. *Math. Oper. Res.* **19** 352–389.

Robinson, S. M. 1979a. Some continuity properties of polyhedral multifunctions. *Math. Programming Study* **14** 206–214.

——. 1979b. Generalized equations and their solution: Part I: Basic theory. *Math. Programming Study* **10** 128–141.

——. 1983. Generalized equations. A. Bachem, M. Grötchel and B. Korte, eds. *Mathematical Programming: The State of the Art, Bonn 1982.* Springer Verlag, Berlin. 346–367.

——. 1992. Normal maps induced by linear transformations. *Math. Oper. Res.* **17** 691–714.

Rockafellar, R. T. 1970. *Convex Analysis.* Princeton University Press, Princeton, N.J.

Rutherford, T. F. 1993. MILES: A mixed inequality and nonlinear equation solver. Working Paper, Department of Economics, University of Colorado, Boulder, CO.

Scholtes, S. 1994. Introduction to piecewise differentiable equations. Preprint No. 53/1994, Institut für Statistik und Mathematische Wirtschaftstheorie, Universität Karlsruhe, Karlsruhe, Germany.

Sun, J. 1986. On Monotropic Piecewise Quadratic Programming. PhD thesis, University of Washington, Seattle, WA.

Todd, M. J. 1976. A note on computing equilibria in economies with activity analysis models of production. *J. Math. Econ.* **6** 135–144.

S. C. Billups: Department of Mathematics, University of Colorado, Denver, Colorado 80217

M. C. Ferris: Computer Sciences Department, University of Wisconsin, Madison, Wisconsin 53706