

GRAPH THEORY, IRREDUCIBILITY, AND STRUCTURAL ANALYSIS OF DIFFERENTIAL-ALGEBRAIC EQUATION SYSTEMS

JOHN D. PRYCE[‡], NEDIALKO S. NEDIALKOV^{†§}, AND GUANGNING TAN^{†¶}

Abstract. The Σ -method for structural analysis of a differential-algebraic equation (DAE) system produces offset vectors from which the sparsity pattern of a system Jacobian is derived. This pattern implies a block-triangular form (BTF) of the DAE that can be exploited to speed up numerical solution.

The paper compares this fine BTF with the usually coarser BTF derived from the sparsity pattern of the signature matrix. It defines a Fine-Block Graph with weighted edges, which gives insight into the relation between coarse and fine blocks, and the permitted ordering of blocks to achieve BTF. It also illuminates the structure of the set of normalised offset vectors of the DAE, e.g. this set is finite if and only if there is just one coarse block.

Key words. differential-algebraic equations, structural analysis, linear assignment problem, block-triangular form, digraph, irreducibility, critical path, PERT

AMS subject classifications. 34A09, 65L80, 41A58, 65F50

1. Introduction. In the paper [17] the authors survey the theory of structural analysis (SA) of differential-algebraic equation systems (DAEs) via the Σ -method [16] based on the DAE's *signature matrix*. In its companion paper [12] they present its implementation in the MATLAB package DAESA. The proof of some results stated there was left to future work. This paper presents those proofs, extends the theory, and introduces terminology and notation to express the relevant ideas concisely.

Understanding DAEs is hard because of the *index* concept: an integer $\nu \geq 0$ (several definitions exist that differ in detail) that measures how “different” a DAE is from an ordinary differential equation, hence how “hard” it inherently is to solve. A DAE can be reduced to standard first-order form $F(t, x, x') = 0$ but this on its own gives no insight into the index. Various methods, e.g., those based on the *derivative array equations* [3], compute the index for a wide class of DAEs but their heavy use of numerical linear algebra makes them expensive.

SA-based methods, being based on the sparsity pattern of the DAE, can be much faster. They can fail, but have a simple numerical test for success. Experience shows they succeed on most DAEs encountered in practice. They provide a systematic solution scheme that can be tailored to various numerical methods. They show exactly what initial values are needed. Till recently the 1989 Pantelides algorithm [14] has been the most used SA method. However, the Σ -method is attractive because of its simplicity; systems up to size 5, say, can be analysed in a few minutes by hand.

Much of SA theory is related to that of sparse systems of linear equations: the notion of a *sparsity pattern*, and of being *structurally nonsingular*; the analysis of dependencies among equations and variables via the system's *bipartite graph* and, when a *transversal* is found, by the corresponding *directed graph*; reduction to *irreducible block triangular form* (BTF), and its relation to *connected components* of these graphs.

[‡]Cardiff School of Mathematics, Cardiff University, UK

[‡]Supported in part by The Leverhulme Trust, UK

[†]Department of Computing and Software, McMaster University, Hamilton, Canada

[§]Supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC)

[¶]Supported in part by the Ontario Research Fund (ORF), Canada

The Σ -method adds to this a *linear assignment problem* (LAP). Being a kind of linear programming problem, this has a dual problem, a solution of which we call an *offset vector*—the name derives from our algorithm to solve a DAE by Taylor series, [10, 11]. Each such vector defines a *System Jacobian matrix*, whose associated sparsity pattern is usually a strict subset of the sparsity pattern of Σ . Thus it gives a *fine-block* irreducible BTF, usually more informative than the *coarse-block* one based on Σ . There is an *essential sparsity pattern*, contained in all the Jacobian patterns but dependent only on Σ , that encapsulates the irreducible fine-block structure.

The main results about these objects were stated in [17] and are proved here. New in this paper is the notion of the *fine-block graph* (FBG), a digraph whose vertices are the fine blocks, with weighted edges. It concisely specifies the set of valid offset vectors. Its strong components can be identified with the coarse blocks. It is close in spirit to the graph used in operational research for *activity analysis* by the Project Evaluation and Review Technique (PERT), and similar terminology is meaningful, such as *critical path* in the graph.

Since this paper says almost nothing about numerics as such, it is really a study of the structure of LAPs, with an emphasis on large sparse ones. We hope the structural insights presented here may find uses in other applications of LAPs. A natural related question is whether they offer a faster way to solve the LAP itself, i.e., find a highest value transversal and dual variables. We are exploring algorithms and aim to report on this in a future paper.

The companion paper [13] studies in detail how BTF can be used (a) to speed up numerical solution, and (b) to find the minimal set of initial values and initial guesses (trial values for a nonlinear solver) that a DAE requires.

Recent work shows that some methods for the DAE initial-value problem (IVP) depend on choosing an offset vector different from the most natural, *canonical*, one. This gives practical interest to analysing the set of valid offset vectors.

The structure of the paper is as follows. Section 2 gives basic results about the LAP that underlies our structural analysis. Section 3 defines permuted square-block forms of a matrix, and when two such forms count as equivalent in our context. Section 4 gives definitions and key results about block-triangular form, BTF. Section 5 studies how the FBG illuminates structural properties of the DAE and its set of offset vectors. Section 6 discusses current, and hopefully future, uses of the theory.

The authors wish to acknowledge Alexei A. Shaleninov, whose work they have recently been made aware of. His 1991 paper [18] pre-dates the first author's work on DAEs starting 1996, and contains a number of results given in [16].

2. Structural analysis and graph theory definitions.

Terms are set in *slanted font* at their defining occurrence.

2.1. Signature matrix, assignment problem, offsets. We consider DAEs of the general form

$$(2.1) \quad f_i(t, \text{the } x_j \text{ and derivatives of them}) = 0, \quad i = 1:n,$$

where¹ the $x_j(t)$, $j = 1:n$ are state variables that are functions of an independent (time) variable t . The f_i can be arbitrary expressions built from the x_j and t using $+$, $-$, \times , \div , other analytic standard functions, and the differentiation operator d^p/dt^p .

¹ $p:q$ for integer p, q denotes either the unordered set, or the enumerated list, of integers i with $p \leq i \leq q$, depending on context.

We call our approach the Σ -method, as it is based on the $n \times n$ *signature matrix* Σ , whose i, j entry σ_{ij} is either an integer ≥ 0 , the order of the highest derivative to which variable x_j occurs in the function f_i ; or $-\infty$ if x_j does not occur in f_i .

The start of the structural analysis is to take Σ as the matrix of a *linear assignment problem* (LAP). The task is to choose a *transversal*, a set T of n positions (i, j) with just one entry in each row and each column, that maximises the sum $\sum_{(i,j) \in T} \sigma_{ij}$, which we call the *value* of T , written $\text{Val}(T)$. That is, we seek a *highest-value transversal*, or HVT—generally nonunique—that gives $\text{Val}(T)$ its largest possible value: this value is denoted $\text{Val}(\Sigma)$.

The DAE is *structurally well-posed* if it has a *finite transversal*, i.e., a T all of whose σ_{ij} are finite (so $\text{Val}(T)$ and hence $\text{Val}(\Sigma)$ are finite); otherwise it is *structurally ill-posed*.

A LAP is a kind of linear programming problem (LPP), so it has a dual LPP. In the formulation we use, this has $2n$ dual variables, $c = (c_1, \dots, c_n)$ and $d = (d_1, \dots, d_n)$, associated with the equations and the variables of (2.1), respectively. The dual² LPP consists of minimising $\sum d_j - \sum c_i$ subject to

$$(2.2) \quad d_j - c_i \geq \sigma_{ij} \text{ for all } i, j.$$

Assume henceforth a structurally well-posed DAE. Then both the primal and the dual LPP have feasible solutions, so the two objective functions have the same, finite, optimal value, see e.g. [7, 9], which is $\text{Val}(\Sigma)$. When the Σ -method succeeds (see §2.2), this equals the *number of degrees of freedom* (DOF) of the DAE: the number of independent initial values required to solve an IVP [16].

Shifting by a constant, i.e. changing to $c'_i = c_i + K$, $d'_j = d_j + K$ for an arbitrary K , does not affect (2.2). Hence various sign constraints can be included in the dual LPP without changing any property that is unaffected by a shift. We impose

$$(2.3) \quad c_i \geq 0 \text{ for all } i,$$

natural in the context of formula (2.7) for the System Jacobian, and of solving DAEs by Taylor series.

LEMMA 2.1. *For any n -vectors c, d , if (2.2) holds—(2.3) is not necessary—and*

$$(2.4) \quad d_j - c_i = \sigma_{ij} \text{ for all } (i, j) \in T$$

holds for some transversal T , then:

(i) *T is a HVT.*

(ii) *Equation (2.4) holds for all HVTs.*

Proof. See [16, Theorem 3.4(iii)], or [10, eqn (2.5)]. \square

Any dual-optimal vectors c, d —equivalently by this Lemma solutions of

$$(2.5) \quad d_j - c_i \geq \sigma_{ij} \text{ with equality on some HVT, hence on all HVTs,}$$

—are termed *general offsets*. If also (2.3) holds they are termed *valid offsets*. Clearly a unique shift K may be chosen to produce *normalised offsets* satisfying

$$(2.6) \quad \min_i c_i = 0.$$

From [16, Theorem 3.6] there exists a unique elementwise smallest solution of (2.3, 2.5), the *canonical offset vector*. It is often convenient to base the SA and numerical solution on these. (Clearly, canonical \implies normalised \implies valid \implies general.)

²Historically, the inequalities (2.2) were identified first as key to the solution process, and the primal problem, the LAP, derived from them.

2.2. System Jacobian. Crucial for a numerical solution method is the $n \times n$ *System Jacobian* matrix $\mathbf{J} = \mathbf{J}(c, d)$ associated with given valid offset vectors c, d . It is defined by

$$(2.7) \quad \mathbf{J}_{ij} = \frac{\partial f_i^{(c_i)}}{\partial x_j^{(d_j)}} = \frac{\partial f_i}{\partial x_j^{(d_j - c_i)}} = \begin{cases} \frac{\partial f_i}{\partial x_j^{(\sigma_{ij})}} & \text{if } d_j - c_i = \sigma_{ij}, \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

for $i, j \in 1:n$. (*Griewank's Lemma*, see e.g. [6], shows the first and second formulae are equivalent.) If a consistent point, in the sense defined in [16], can be found at which \mathbf{J} is nonsingular, then at least locally there exists a solution path of the DAE through that point. In this case the Σ -method is said to *succeed*. Various numerical methods based on the SA can then be used to compute such a path, such as:

- Taylor series expansion [11];
- index reduction using dummy derivatives [8] followed by solution with a standard index-1 DAE solver or implicit ODE solver.

Example 2.1. Below is shown the DAE of the simple pendulum in Cartesian coordinates, together with its signature matrix and System Jacobian, using canonical offsets.

$$(2.8) \quad \begin{array}{l} \text{PEND} \\ 0 = A = x'' + x\lambda \\ 0 = B = y'' + y\lambda - G \\ 0 = C = x^2 + y^2 - L^2 \end{array}, \quad \Sigma = \begin{array}{c} x \quad y \quad \lambda \quad c_i \\ A \begin{bmatrix} 2^\bullet & & & \\ & 2^\circ & & \\ & & 0^\bullet & \\ & & & 2 \end{bmatrix} \\ B \begin{bmatrix} & & & \\ & & & \\ & & & \\ & & & 2 \end{bmatrix} \\ C \begin{bmatrix} & & & \\ & & & \\ & & & \\ & & & 2 \end{bmatrix} \\ d_j \quad 2 \quad 2 \quad 0 \end{array}, \quad \mathbf{J} = \begin{array}{c} x'' \quad y'' \quad \lambda \\ A \begin{bmatrix} 1 & 0 & x \\ 0 & 1 & x \\ 2x & 2y & 0 \end{bmatrix} \\ B \\ C'' \end{array}$$

Here x, y, λ are state variables; G, L are constants > 0 . A blank in Σ means $-\infty$.

The rows and columns of Σ are annotated on the left and top with the names of the functions A, B, C and the variables x, y, λ , and on the right and bottom with valid (in this case canonical) offsets. There are two HVTs, marked with \bullet and \circ .

An alternative to showing the offsets explicitly is to show the i th function [respectively j th variable] differentiated to order c_i [resp. d_j], as shown for \mathbf{J} . It is a reminder of the first formula for \mathbf{J} in (2.7): here, $\mathbf{J} = \partial(A, B, C'')/\partial(x'', y'', \lambda)$. \square

2.3. Graph theory notation and terminology. For clarity we state the notions from graph theory that we use, see for instance [1, 4, 5, 15]. A *directed graph*, or *digraph*, is a pair $G = (V, E)$ where V is a nonempty set whose elements are the *vertices* (or *nodes*) and E is a set of ordered pairs $e = (u, v)$ of vertices. e is the edge *from* u *to* v , often written $u \rightarrow v$. By the nature of a set there is at most one edge from u to v . We require $u \neq v$: there are no edges from a vertex to itself.

A *path* in G is a sequence of edges where the end of each one is the start of the next. It may be written $v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_m$, where $v_{i-1} \rightarrow v_i$ is an edge for each $i = 1:m$. It is allowed to have repeated vertices and/or edges. If there is a path from u to v (i.e., a path having $v_0 = u$ and $v_m = v$) then v is *reachable* from u .

A *cycle*, or closed path, is a path of nonzero length for which $v_0 = v_m$. It may have repeated vertices and/or edges in addition to this.

A general *subgraph* of G means (V', E') where V' and E' are arbitrary subsets of V and E . The subgraph *generated* by V' has $E' = \{\text{all } e \in E \text{ joining elements of } V'\}$.

A digraph G is *acyclic* if it has no cycles. It is *strongly connected* if for any two vertices, each is reachable from the other.

An *undirected graph* is one whose edges are bidirectional (equivalently, they come in pairs $u \rightarrow v$ and $v \rightarrow u$). For such a graph, if v is reachable from u then u is reachable from v , and one just calls it *connected* rather than strongly connected.

Every digraph decomposes into *strong components*. Namely there is a unique partition $\{V_1, \dots, V_q\}$ of V such that the subgraph G_i generated by V_i —a strong component—is strongly connected for each i , and every cycle in G is contained in one of the components. Thus given two strong components, one may be reachable from the other in G , but never in both directions.

3. Permuted square-block forms. This section defines block form to be the result of partitioning the equations and variables into matching sets that may be regarded as sub-DAEs of the whole system. Suitably permuted, they define a block structure on relevant matrices, such that blocks on the main diagonal are square. It introduces the notion of *emblem* to express when two such forms count as “the same”.

3.1. Definitions. A permuted form of a DAE (2.1) is created by forming permutations $\tilde{f} = (\tilde{f}_1, \dots, \tilde{f}_n)$ and $\tilde{x} = (\tilde{x}_1, \dots, \tilde{x}_n)$ of the vectors of equations f_i and variables x_j . A (permuted) *square-block form* is a triple $\mathcal{B} = (\tilde{f}, \tilde{x}, N)$ where \tilde{f} and \tilde{x} are as above, and N is a vector (N_1, \dots, N_p) of positive integers summing to n , which put \tilde{f} and \tilde{x} in block form (F_1, \dots, F_p) and (X_1, \dots, X_p) , where each of sub-vectors F_l and X_l has N_l elements for $l = 1:p$.

Then the signature matrix Σ and a Jacobian \mathbf{J} are permuted correspondingly and put in $p \times p$ block form:

$$(3.1) \quad \tilde{\Sigma} = \begin{bmatrix} \tilde{\Sigma}_{11} & \cdots & \tilde{\Sigma}_{1p} \\ \vdots & & \vdots \\ \tilde{\Sigma}_{p1} & \cdots & \tilde{\Sigma}_{pp} \end{bmatrix}, \quad \tilde{\mathbf{J}} = \begin{bmatrix} \tilde{\mathbf{J}}_{11} & \cdots & \tilde{\mathbf{J}}_{1p} \\ \vdots & & \vdots \\ \tilde{\mathbf{J}}_{p1} & \cdots & \tilde{\mathbf{J}}_{pp} \end{bmatrix},$$

where the sub-matrix in position (k, l) has size N_k by N_l . In particular each diagonal sub-matrix $\tilde{\Sigma}_{ll}$ or $\tilde{\mathbf{J}}_{ll}$ is square of size N_l .

The equations (rows) and variables (columns) have names that form labels attached to them and are carried with them under permutations. In the example DAEs in this paper, specific names are given: e.g. (2.8) has A, B, C for the equations and x, y, λ for the variables, and these are the labels. For a generic DAE, we name the functions f_1, \dots, f_n , the variables x_1, \dots, x_n . As labels we use either these names, or the indices written boxed—e.g. $\boxed{2}$ for f_2 or x_2 .

In examples, horizontal and vertical lines may display block structure.

In the context of square-block form a *block* means, not a sub-matrix such as $\tilde{\Sigma}_{kl}$, but a set of N_l equations and corresponding set of N_l variables, associated with the l th diagonal sub-matrix. Depending on context (see §3.2) it may be treated as a pair of sets (unordered) of row/column labels; or a pair of lists (ordered) of such labels; or in other ways.

For instance, the following versions of a signature matrix, annotated with function- and variable-labels, describe the same permutation of (2.8) and an accompanying square-block form. The first uses the given names A, B, C, x, y, λ ; the second uses the generic indexing.

The permuted offsets are also shown. It is easy to see from their definition (2.2) that they can be regarded as attached to rows and columns and carried around like the equation- and variable-labels. That is, permuted \tilde{c}, \tilde{d} are general [or valid or

normalised or canonical] offsets of the permuted problem iff c, d are general [or etc.] offsets of the original problem.

$$(3.2) \quad \tilde{\Sigma} = \begin{array}{c} C \\ A \\ B \\ \tilde{d}_j \end{array} \left[\begin{array}{cc|c} y & \lambda & x \\ 0 & & 0 \\ & 0 & 2 \\ 2 & 0 & \end{array} \right] \begin{array}{c} \tilde{c}_i \\ 2 \\ 0 \\ 0 \end{array} = \begin{array}{c} \boxed{3} \\ \boxed{1} \\ \boxed{2} \\ \tilde{d}_j \end{array} \left[\begin{array}{cc|c} \boxed{2} & \boxed{3} & \boxed{1} \\ 0 & & 0 \\ & 0 & 2 \\ 2 & 0 & \end{array} \right] \begin{array}{c} \tilde{c}_i \\ 2 \\ 0 \\ 0 \end{array}.$$

The following describes (3.2) in generic block form:

$$\tilde{\Sigma} = \begin{array}{c} F_1 \\ F_2 \end{array} \left[\begin{array}{c|c} X_1^T & X_2^T \\ \hline \tilde{\Sigma}_{11} & \tilde{\Sigma}_{12} \\ \tilde{\Sigma}_{21} & \tilde{\Sigma}_{22} \end{array} \right], \quad \text{where } F_1 = \begin{pmatrix} C \\ A \end{pmatrix}, F_2 = (B), X_1 = \begin{pmatrix} y \\ \lambda \end{pmatrix}, X_2 = (x).$$

The relation between original and permuted rows and columns, and between original and permuted offset vectors, is

$$\begin{aligned} \tilde{f}_i &= f_{\rho(i)}, & \tilde{x}_j &= x_{\kappa(j)}, \\ \tilde{c}_i &= c_{\rho(i)}, & \tilde{d}_j &= d_{\kappa(j)}, \end{aligned}$$

where $\rho(i)$ and $\kappa(j)$ are the labels of row i and column j in the permuted form for $i, j = 1 : n$. Clearly, regarded as maps of $1 : n$ to itself, ρ and κ are permutations. For example in (3.2) one has $\rho(1) = 3, \rho(2) = 1, \rho(3) = 2$.

Similarly, for a signature matrix and Jacobian the relation between original Σ, \mathbf{J} and permuted $\tilde{\Sigma}, \tilde{\mathbf{J}}$ is

$$(3.3) \quad \tilde{\sigma}_{ij} = \sigma_{\rho(i), \kappa(j)}, \quad \tilde{\mathbf{J}}_{ij} = \mathbf{J}_{\rho(i), \kappa(j)},$$

and for a pattern A and permuted pattern \tilde{A} regarded as sets—including the case that A is a transversal and \tilde{A} the permuted transversal—

$$(3.4) \quad (i, j) \in \tilde{A} \iff (\rho(i), \kappa(j)) \in A,$$

For $l \in 1 : p$ define

$$(3.5) \quad B_l = \text{the set of integer indices } i \text{ that belong to block } l,$$

that is,

$$B_1 = 1 : N_1, \quad B_2 = (N_1+1) : (N_1+N_2), \quad \dots,$$

forming a partition of $1 : n$. Note i denotes a current row or column *index*, not an original row or column *label*, so the notation applies equally to rows and columns.

We also use $\text{blockOf}(i)$ to denote the l such that $i \in B_l$, so for $i \in 1 : n, l \in 1 : p$

$$(3.6) \quad \begin{aligned} \text{blockOf}(i) = l &\iff i \in B_l \\ &\iff \sum_{r=1}^{l-1} N_r < i \leq \sum_{r=1}^l N_r. \end{aligned}$$

To know which block an original *label* belongs to, use the inverse permutations to ρ and κ . For example, original equation \boxed{i} lies in block $l = \text{blockOf}(\rho^{-1}(i))$.

3.2. Equivalence and the emblem. When comparing different permuted square-block forms of a given DAE, one must specify when a block in one is considered “the same” as a block in another; hence, when two block forms are considered “the same”. We count a block as unchanged by a permutation of equations and/or variables purely within that block; and the whole block form as unchanged by a permutation of the blocks as a whole. For instance in the forms

$$(3.7) \quad \begin{array}{c} \begin{array}{c} \begin{array}{c} w \quad x \quad y \quad z \\ \hline \begin{array}{c} f \\ g \\ h \\ k \end{array} \begin{array}{c} \left[\begin{array}{c|c|c|c} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{array} \right] \end{array} \end{array} \end{array}, \quad \begin{array}{c} \begin{array}{c} x \quad w \quad y \quad z \\ \hline \begin{array}{c} g \\ f \\ k \\ h \end{array} \begin{array}{c} \left[\begin{array}{c|c|c|c} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{array} \right] \end{array} \end{array} \end{array}, \\ \begin{array}{c} \begin{array}{c} x \quad w \quad y \quad z \\ \hline \begin{array}{c} f \\ g \\ k \\ h \end{array} \begin{array}{c} \left[\begin{array}{c|c|c|c} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{array} \right] \end{array} \end{array} \end{array}, \quad \begin{array}{c} \begin{array}{c} z \quad y \quad w \quad x \\ \hline \begin{array}{c} h \\ k \\ f \\ g \end{array} \begin{array}{c} \left[\begin{array}{c|c|c|c} \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \\ \times & \times & \times & \times \end{array} \right] \end{array} \end{array} \end{array}, \end{array}$$

—assuming the matrix bodies are permuted to match—the block with equations h, k and variables y, z is “the same” in all of them. Forms 1 and 2 are “the same” as a whole because both have f matched with w ; g with x ; and h, k with y, z —though in different orders. Form 3 has block sizes the same as forms 1 and 2 but differs from them because it has f matched with x and g with w . Form 4 must be different from forms 1, 2 and 3, since it has different block sizes.

We describe this formally as follows. In a square-block form \mathcal{B} , the *row-set* R_l of block l is the set of equation-labels of that block; its *column-set* C_l is its set of variable-labels. By definition, each R_l is paired with a C_l having the same number of elements: $|R_l| = |C_l| = N_l$. Its *emblem* is the set of pairs

$$(3.8) \quad \mathcal{E} = \text{emb}(\mathcal{B}) = \{(R_1, C_1), \dots, (R_p, C_p)\}.$$

Abstractly, an emblem is *an unordered set of ordered pairs of unordered sets* as shown in (3.8), where $\{R_1, \dots, R_p\}$ and $\{C_1, \dots, C_p\}$ are partitions³ of the sets of equation- and variable-labels respectively—of $1:n$ if using generic labeling. Its purpose is to specify a square-block structure up to equivalence.

When a permuted block form is shown graphically, the main block-diagonal indicates the (R_l, C_l) pairing. In (3.7), forms 1 and 2 are equivalent, having emblem

$$(3.9) \quad \mathcal{E}_1 = \{(\{f\}, \{w\}), (\{g\}, \{x\}), (\{h, k\}, \{y, z\})\},$$

while forms 3 and 4 have the different emblems

$$\begin{aligned} \mathcal{E}_2 &= \{(\{f\}, \{x\}), (\{g\}, \{w\}), (\{h, k\}, \{y, z\})\} \quad \text{and} \\ \mathcal{E}_3 &= \{(\{f, g\}, \{w, x\}), (\{h, k\}, \{y, z\})\}. \end{aligned}$$

A given permuted block form *enumerates* its emblem, which means listing the pairs (R_l, C_l) , and the members of each R_l and C_l , in some order—using square

³A partition of a set X is a pairwise disjoint family of nonempty subsets of X whose union is X .

bracket notation for lists, this amounts to “turning braces to brackets” in an explicit description. For instance, forms 1 and 2 of (3.7) enumerate \mathcal{E}_1 of (3.9) in two ways:

$$\begin{aligned} & [([f], [w]), ([g], [x]), ([h, k], [y, z])], \quad \text{and} \\ & [([g], [x]), ([f], [w]), ([k, h], [y, z])]. \end{aligned}$$

From the above definitions we clearly have:

THEOREM 3.1. *Two square-block forms have the same emblem, and are then called block-equivalent, if and only if one can be converted to the other by permuting the blocks as a whole, and/or equations or variables within an individual block.*

This is purely about block structure—assigning equations and variables to blocks. It says nothing about matrix values σ_{ij} , etc. or, until Section 4, about sparsity issues.

In terms of the DAE, we call blocks *parts*. Part l means the system of equations and variables specified by R_l and C_l , regarded as a DAE of size N_l in its own right, with the relevant sub-matrices of (permuted) Σ and \mathbf{J} , sub-vectors of offsets, etc. In Σ and \mathbf{J} the off-diagonal blocks in a part’s block-column [-row] indicate how it influences [is influenced by] other parts.

In terms of an underlying physical model, a part is usually an identifiable component of the whole system, though this is not mathematically necessary. E.g., (3.2) shows a split of the pendulum system into parts with no obvious physical meaning.

4. Block-triangular form within the DAE. In this section, §4.1 defines the relevant sparsity patterns and proves some basic BTF facts. §4.2 introduces irreducibility, and coarse and fine block form, and proves the key result that the fine-block form of a DAE is unique, in the sense that its emblem is independent of the offsets used to compute it. §4.3 notes some caveats regarding this analysis.

4.1. Sparsity patterns and BTF.

4.1.1. Sparsity pattern definitions. A BTF of (2.1) is obtained from a *sparsity pattern*, which is some subset A of $\{1, \dots, n\}^2$, the $n \times n$ matrix positions (i, j) for i and j from 1 to n . When appropriate, we identify A with its $n \times n$ *incidence matrix*

$$(a_{ij}) \text{ where } a_{ij} = 1 \text{ if } (i, j) \in A, 0 \text{ otherwise.}$$

In the context of a permuted pattern \tilde{A} , the (i, j) always refer to the *original* equation- and variable-indices (the labels). For example, the set A where the $\sigma_{ij} > -\infty$ for the pendulum DAE (2.8), in the permuted form of (3.2), may be shown:

$$\tilde{A} = \begin{array}{c} C \\ A \\ B \end{array} \begin{array}{c|cc} & y & \lambda & x \\ \hline & 1 & & 1 \\ & & 1 & 1 \\ \hline & 1 & 1 & \end{array} = \begin{array}{c} \boxed{3} \\ \boxed{1} \\ \boxed{2} \end{array} \begin{array}{c|cc} \boxed{2} & \boxed{3} & \boxed{1} \\ \hline & 1 & 1 & 1 \\ \hline & 1 & 1 & \end{array}.$$

It is written as a sparse incidence matrix, i.e. blank entries denote zero. The element in row 3, column 1 is the (B, y) entry, or generic notation the (f_2, x_2) entry (since $\tilde{f}_3 = B = f_2$ and $\tilde{x}_1 = y = x_2$ for this permutation).

A natural sparsity pattern for a DAE is the set where the entries of Σ are finite:

$$(4.1) \quad S = \{ (i, j) \mid \sigma_{ij} > -\infty \} \quad (\text{the sparsity pattern of } \Sigma).$$

However, a more informative BTF comes from the sparsity pattern of the Jacobian \mathbf{J} defined in (2.7) (but see the caveats in §4.3). It depends, as does \mathbf{J} , on the (valid) offsets c, d used:

$$(4.2) \quad S_0 = S_0(c, d) = \{(i, j) \mid d_j - c_i = \sigma_{ij}\} \quad (\text{the sparsity pattern of } \mathbf{J}).$$

An important, less obvious set⁴, is

$$(4.3) \quad S_{\text{ess}} = \text{the union of all HVTs of } \Sigma. \quad (\text{the essential sparsity pattern}).$$

Since $d_j - c_i = \sigma_{ij}$ holds on each HVT by (2.5), and implies $\sigma_{ij} > -\infty$, we have

$$S_{\text{ess}} \subseteq S_0(c, d) \subseteq S \quad \text{for any } c, d.$$

Experience suggests that in applications, a BTF based on S_0 is usually significantly finer than one based on S . We refer to any instance of the former as a *fine* BTF, and to the latter as the *coarse* BTF.

4.1.2. Definition of BTF. Following common usage in the literature, e.g. [15], also used by DAESA, we use *upper* block-triangular form as below:

$$\Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} & \cdots & \Sigma_{1p} \\ -\infty & \Sigma_{22} & \cdots & \Sigma_{2p} \\ \vdots & \ddots & \ddots & \vdots \\ -\infty & \cdots & -\infty & \Sigma_{pp} \end{bmatrix}, \quad \mathbf{J} = \begin{bmatrix} \mathbf{J}_{11} & \mathbf{J}_{12} & \cdots & \mathbf{J}_{1p} \\ 0 & \mathbf{J}_{22} & \cdots & \mathbf{J}_{2p} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \mathbf{J}_{pp} \end{bmatrix}, \quad A = \begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1p} \\ \emptyset & A_{22} & \cdots & A_{2p} \\ \vdots & \ddots & \ddots & \vdots \\ \emptyset & \cdots & \emptyset & A_{pp} \end{bmatrix}.$$

That is, given a permuted square-block form—an enumeration of some emblem—the permuted signature matrix Σ is in BTF if the below-diagonal sub-matrices, Σ_{kl} where $k > l$, are all $-\infty$. A corresponding Jacobian \mathbf{J} is in BTF if the below-diagonal sub-matrices are all zero. When a sparsity pattern A is viewed as a set of (i, j) 's, its sub-matrices are the pairwise disjoint *sub-patterns*, the sets

$$(4.4) \quad A_{kl} = A \cap (B_k \times B_l),$$

(where the B s are as in (3.5)) and A is in BTF if the below-diagonal A_{kl} are empty.

Clearly, an arbitrary permutation of rows or columns within a block—an enumeration of the emblem that keeps the same pairs (R_l, C_l) , in the same order, but lists the members of some R_l or C_l differently—preserves BTF, in the sense that the result is still in upper BTF with the same *sequence* of block sizes N_1, N_2, \dots .

A permutation of the blocks as a whole—an enumeration that lists the (R_l, C_l) in a different order—produces a new block form that is usually not in BTF (unless, for instance, the original form is actually block-diagonal).

Basing BTF on a sparsity pattern S_0 of (4.2) puts the corresponding \mathbf{J} into BTF by definition but may not do so for Σ , as the next example shows.

Example 4.1. Suppose in the pendulum DAE (2.8), x is changed to x' in the third equation. This gives the following. The canonical offsets, shown on Σ , are used to compute \mathbf{J} , and indicated on \mathbf{J} using the convention described in Example 2.1.

$$\begin{array}{l} 0 = A = x'' + x\lambda \\ 0 = B = y'' + y\lambda - G \\ 0 = C = (x')^2 + y^2 - L^2 \end{array} \quad \Sigma = \begin{array}{c} A \\ B \\ C \\ d_j \end{array} \begin{array}{ccc} x & y & \lambda \\ \left[\begin{array}{ccc} 2 & & 0^\circ \\ & 2^\circ & 0 \\ 1^\circ & 0 & \end{array} \right] & & \\ & 2 & 2 & 0 \end{array} \quad \mathbf{J} = \begin{array}{c} A \\ B \\ C' \end{array} \begin{array}{ccc} x'' & y'' & \lambda \\ \left[\begin{array}{ccc} 1 & 0 & x \\ 0 & 1 & y \\ 2x' & 0 & 0 \end{array} \right] & & \end{array}$$

⁴It was called S_{00} in [16]; see also [10].

Note σ_{31} changes from 0 to 1, c_3 from 2 to 1, and \mathbf{J}_{32} is now 0. Hence

$$S_0(c, d) = \begin{bmatrix} 1 & & 1 \\ & 1 & 1 \\ 1 & & \end{bmatrix} \text{ is a proper subset of } S = \begin{bmatrix} 1 & & 1 \\ & 1 & 1 \\ 1 & 1 & \end{bmatrix}.$$

The following form puts \mathbf{J} into BTF with three blocks of size 1:

$$\tilde{\Sigma} = \begin{array}{c} B \\ A \\ C \\ d_j \end{array} \begin{array}{ccc} y & \lambda & x \\ \left[\begin{array}{ccc} 2^\circ & 0 & \\ & 0^\circ & 2 \\ 0 & & 1^\circ \end{array} \right] & & \\ 2 & 0 & 2 \end{array} \begin{array}{c} c_i \\ 0 \\ 0 \\ 1 \end{array} \quad \tilde{\mathbf{J}} = \begin{array}{c} B \\ A \\ C' \end{array} \begin{array}{ccc} y'' & \lambda & x'' \\ \left[\begin{array}{c|c|c} 1 & y & 0 \\ 0 & x & 1 \\ 0 & 0 & 2x' \end{array} \right] \end{array}$$

but does not put Σ into BTF. \square

4.1.3. Results on transversals and BTF. Several results of this section apply to an arbitrary pattern A , but, to simplify notation, the proofs “assume without loss” that A is already in BTF. For instance, Theorem 4.1 is to be interpreted as follows.

T is a transversal of A , and \tilde{A} is a permuted form of A that is in BTF. Then \tilde{T} , the corresponding permutation of T following (3.4), which is a transversal of \tilde{A} , is contained in the union of the diagonal blocks of \tilde{A} .

THEOREM 4.1. *Any transversal of (i.e., contained in) a sparsity pattern A is contained in the union of the diagonal blocks of any BTF of A .*

Proof. As said above, we may assume that A is already in BTF. The proof is by induction on p , the number of blocks. The result is trivially true in the base case $p = 1$.

Assume inductively the result holds for a BTF with $p - 1$ blocks where $p > 1$. Let A , of size n , be in upper BTF with p blocks, and T be a transversal of it. With the notation above, its last, p th, block has size N_p . By the definition of BTF, A has no elements to the left of the last block, that is in the last N_p rows and first $n - N_p$ columns. Hence the N_p elements of T in these rows are contained in the last block, and must therefore occupy all N_p columns of the last block.

Thus T' , comprising the remaining $n - N_p$ elements of T , occupies the first $n - N_p$ rows and columns and therefore is a transversal of A' , the intersection of A with these rows and columns. The first $p - 1$ blocks of the given BTF form a BTF of A' . By the inductive hypothesis T' is contained in the diagonal blocks of this BTF. This and the previous paragraph show T is contained in the diagonal blocks of the whole BTF, proving the inductive step and completing the proof. \square

LEMMA 4.2.

- (i) *Any transversal of a Jacobian sparsity pattern $S_0 = S_0(c, d)$, see (4.2), is a HVT of Σ .*
- (ii) *For any BTF of a Jacobian sparsity pattern S_0 , the restriction to a block of any valid offset vector of Σ (not necessarily the one from which S_0 is derived) is a valid offset vector of the corresponding diagonal block of Σ .*

Proof. Part (i) is a restatement of Lemma 2.1 (i).

For part (ii), again we may assume S_0 is already in BTF. Let the blocks of S_0 and Σ be $(S_0)_{kl}$ and Σ_{kl} for $k \in 1:p$, $l \in 1:p$. Let (c, d) be a valid offset vector of Σ , and write it in block form $((\mathbf{c}_1, \dots, \mathbf{c}_p), (\mathbf{d}_1, \dots, \mathbf{d}_p))$ where each of the sub-vectors \mathbf{c}_l and \mathbf{d}_l has length N_l , following the block form of Σ . Fix $l \in 1:p$. It is required to

show $(\mathbf{c}_l, \mathbf{d}_l)$ is a valid offset vector of block Σ_{ll} , regarded as a signature matrix in its own right.

Write $\Sigma_{ll}, (S_0)_{ll}, \mathbf{c}_l, \mathbf{d}_l$ as $\Sigma^*, S_0^*, c^*, d^*$ to simplify index notation. Now $d_j - c_i = \sigma_{ij}$ holds on some HVT of Σ . This HVT is a transversal of S_0 so by Theorem 4.1 it is contained in the union of the diagonal blocks of S_0 . Hence its intersection with S_0^* is a transversal of S_0^* , a fortiori of Σ^* . Hence $d_j^* - c_i^* = \sigma_{ij}^*$ holds, for the offsets restricted to Σ^* , on a transversal of Σ^* . Clearly also $d_j^* - c_i^* \geq \sigma_{ij}^*$ and $c_i^* \geq 0$ hold. Thus c^*, d^* satisfy the conditions (2.5) for being a valid offset vector of Σ^* . \square

4.2. Irreducibility. This subsection summarises some classic graph theory definitions and results, see e.g. [1, 4, 5, 15].

A sparsity pattern A is *structurally nonsingular* if it contains at least one transversal. A structurally nonsingular A is *irreducible*, if it cannot be permuted to a non-trivial BTF (one with $p > 1$ blocks); otherwise it is reducible.

The *bipartite graph* of A is the undirected graph whose $2n$ vertices are labels for the n rows and n columns, and which has an edge between row i and column j whenever $(i, j) \in A$. In this context a transversal is also called a *maximal matching*.

If A is structurally nonsingular, one can choose (arbitrarily) some transversal T to be *distinguished*, thus specifying a row–column matching: row i matches column j if $(i, j) \in T$. For convenience, assume rows or columns or both are permuted to put T on the main diagonal, so that row i matches column i . The *graph of A relative to T* , which we denote by $\text{Graph}(A, T)$, is then the digraph whose vertices are the indices $1:n$ and which has an edge from i to j if $(i, j) \in A$.

THEOREM 4.3.

- (a) *The following are equivalent:*
- (i) *A is structurally nonsingular.*
 - (ii) *A has the Hall Property: for $r = 1, \dots, n$, any set of r columns of A contains elements of at least r rows.*
- (b) *The following are equivalent:*
- (i) *A is irreducible.*
 - (ii) *A has the Strong Hall Property: for $r = 1, \dots, n - 1$, any set of r columns of A contains elements of at least $r + 1$ rows.*
 - (iii) *Every element of A is on a transversal, and the bipartite graph of A is connected.*
 - (iv) *For any transversal T , the digraph $G = \text{Graph}(A, T)$ is strongly connected.*

If any diagonal block A_{ll} of a BTF of A is reducible (as a sparsity pattern in its own right), this lets one refine the whole BTF of A , splitting the l th block into two or more smaller blocks. Repeating this till no further refinement is possible, we see there exists an *irreducible BTF* of A for which each diagonal block is irreducible. In fact, see [5], this is unique up to reordering of rows and columns within a block, and possible reordering of the blocks—both of which leave the emblem (see (3.8)) unchanged. Equivalently said,

THEOREM 4.4. *Each structurally nonsingular pattern A has a unique irreducible emblem, denoted $\text{emb}(A)$, which is the emblem of every irreducible BTF of A .*

The following key result ties any *Jacobian* sparsity pattern to the *essential* sparsity pattern.

THEOREM 4.5. *For a given signature matrix Σ , let (c, d) be a valid offset vector, let $S_0 = S_0(c, d)$ be its Jacobian sparsity pattern (4.2). Let a permuted square-block form be chosen that puts S_0 into irreducible BTF.*

Then, modulo the permutation, the essential sparsity pattern S_{ess} of Σ , see (4.3), equals the union of the diagonal blocks. Hence S_0 and S_{ess} have the same irreducible emblem.

Proof. That is, let a $\tilde{}$ denote the permuted versions of S_0, S_{ess} , etc. and let (\tilde{S}_{kl}) be the block decomposition of \tilde{S}_0 :

$$\tilde{S}_{kl} = \tilde{S}_0 \cap (B_k \times B_l),$$

in the notation of (3.5). It is required to show

$$(4.5) \quad \tilde{S}_{\text{ess}} = \tilde{S}_{11} \cup \dots \cup \tilde{S}_{pp}.$$

As usual, we assume the matrices are already permuted, and drop the $\tilde{}$ superscripts.

Let the union on the right of (4.5) be called S_{diag} . That $S_{\text{ess}} \subseteq S_{\text{diag}}$ is immediate from Theorem 4.1. For the reverse, take any $(i, j) \in S_{\text{diag}}$. Then (i, j) is in some diagonal block S_{qq} , which is itself irreducible by the definition of irreducible BTF. By Theorem 4.3 (b)(iii), (i, j) is on some transversal of S_{qq} . Taking the union of this with arbitrary transversals (which must exist) of the remaining blocks gives a transversal T of S_0 . By Lemma 4.2 (i), T is a HVT of Σ . Thus $(i, j) \in T \subseteq S_{\text{ess}}$. This shows $S_{\text{diag}} \subseteq S_{\text{ess}}$, completing the proof of (4.5).

For the last assertion of the theorem, by construction the emblem \mathcal{E} of the chosen BTF is the irreducible emblem of S_0 . By (4.5) it puts S_{ess} into block-triangular—indeed block-diagonal—form. Since each S_{ll} is irreducible in its own right, this form is an irreducible BTF of S_{ess} , so the irreducible emblem (defined in Theorem 4.4) of S_{ess} is also \mathcal{E} , as required. \square

COROLLARY 4.6.

- (i) Any irreducible BTF of S_{ess} is block-diagonal, not merely block-triangular.
- (ii) All Jacobian sparsity patterns $S_0 = S_0(c, d)$ of Σ have the same irreducible emblem, which equals that of S_{ess} .

Proof. Immediate from the proof of the preceding theorem. \square

Since this emblem depends only on the signature matrix Σ we call it *the irreducible emblem of Σ* , or of the underlying DAE, and denote it by

$$(4.6) \quad \text{emb}(\Sigma).$$

Any square-block form having this emblem is called an *irreducible fine-block form* of the DAE or an associated matrix (or pattern)—whether or not it puts the matrix or pattern into BTF.

Thus, up to possible reordering, all Jacobian sparsity patterns $S_0(c, d)$ of Σ have the same block sizes N_1, \dots, N_p , and identical diagonal blocks, in their irreducible BTFs. We call these the *fine blocks* of the DAE.

The blocks of an irreducible BTF of S , the sparsity pattern of Σ , are the *coarse blocks*. A signature matrix or DAE is *fine-block* [respectively *coarse-block*] *irreducible* if it has only one fine [coarse] block.

Depending on the DAE, some ordering of blocks in the irreducible BTF of a Jacobian sparsity pattern may be arbitrary; some may be dictated by a specific choice of offsets c and d ; some may be inherent in the DAE. This issue is clarified by the fine-block graph introduced in Section 5.

THEOREM 4.7.

- (i) If there exist valid offsets c, d whose Jacobian pattern $S_0(c, d)$ is irreducible, then the DAE is fine-block irreducible. Conversely if the DAE is fine-block irreducible then $S_0(c, d)$ is irreducible for all valid offset vectors c, d .
- (ii) If the DAE is fine-block irreducible then valid offsets c, d are uniquely determined up to a constant. That is, $c_i = \widehat{c}_i + K$, $d_i = \widehat{d}_i + K$ for some $K \geq 0$, where $(\widehat{c}, \widehat{d})$ is the canonical offset vector.

Equivalently said: if the DAE is fine-block irreducible then its only normalised (see (2.6)) offset vector is the canonical one.

Proof. (i) Immediate from Corollary 4.6 (ii).

(ii) We use the bipartite graph characterisation in Theorem 4.3. Suppose the DAE is fine-block irreducible and let valid offsets c, d be given. Let \mathcal{G} be the bipartite graph of $S_0(c, d)$ and denote its row and column vertices R_1, \dots, R_n and C_1, \dots, C_n respectively. Choose some arbitrary $i_0 \in 1:n$; for convenience let it be some i where the canonical offsets \widehat{c}, \widehat{d} have $\widehat{c}_i = 0$. Let c, d be any other valid offset vector and define $K = c_{i_0}$.

Take an arbitrary $i \in 1:n$. By Theorem 4.3 (b)(iii) there is a path in \mathcal{G} from R_{i_0} to R_i , of the form

$$R_{i_0} \rightarrow C_{j_1} \rightarrow R_{i_1} \rightarrow \dots \rightarrow C_{j_m} \rightarrow R_{i_m} = R_i$$

($2m$ edges in total). Otherwise said, both (i_{r-1}, j_r) and (i_r, j_r) are in $S_0(c, d)$ for $r = 1:m$. By the cited theorem, every element of $S_0(c, d)$ is on a transversal, so we have $2m$ equalities

$$d_{j_1} - c_{i_0} = \sigma_{i_0, j_1}, \quad d_{j_1} - c_{i_1} = \sigma_{i_1, j_1}, \quad \dots$$

Subtracting these in pairs we have

$$\begin{aligned} c_{i_1} - c_{i_0} &= \sigma_{i_0, j_1} - \sigma_{i_1, j_1} &&= s_1, \text{ say;} \\ c_{i_2} - c_{i_1} &= \sigma_{i_1, j_2} - \sigma_{i_2, j_2} &&= s_2; \\ \dots & \dots && \end{aligned}$$

Adding these from 1 to m gives

$$c_i - K = c_{i_m} - c_{i_0} = s_1 + \dots + s_m,$$

where the RHS is independent of the particular offsets c, d chosen. In particular taking c, d to be \widehat{c}, \widehat{d} (for which $K = 0$) we see the RHS equals \widehat{c}_i . Thus we have shown there is K (independent of i) such that $c_i = \widehat{c}_i + K$ for each $i \in 1:n$, as required.

If $K \neq 0$ then either $K < 0$ and $c_i < 0$ for some i so c is not valid; or $K > 0$ and $c_i > 0$ for all i so c is not normalised. This proves the final assertion. \square

4.3. Caveats. Some comments are needed about sparsity analysis. First, if the DAE is structurally well-posed (§2.1), an analysis based on S_0 can always be carried out. However, it is only meaningful if the method is known to succeed: that is, if a consistent point where \mathbf{J} is nonsingular has been found. Second, S_0 is the set of positions where $\mathbf{J}_{i,j}$ is nonzero (a) structurally and (b) as far as DAESA can tell from the equations as written.

As regards (a), \mathbf{J} varies along the solution, and a generically nonzero $\mathbf{J}_{i,j}$ may vanish at certain t values. For (b), algebraic cancellation may make some $\mathbf{J}_{i,j}$ identically zero for reasons that DAESA does not detect. For instance, suppose equation f_1

is $(tx_1)' - tx_1' + x_2 = 0$, which simplifies to $x_1 + x_2 = 0$. DAESA applies quasilinear-ity analysis based on an operator-overloading computation and does not do symbolic simplification, so it finds $\sigma_{11} = 1$ instead of the true $\sigma_{11} = 0$. Such overestimation of elements of Σ looks dangerous, but—provided a nonsingular \mathbf{J} is in due course found—can do no worse than cause some zero entries of \mathbf{J} to be treated as nonzero, and make numerical solution slightly less efficient than it could be: see [10, §5].

5. The set of normalised offsets and the fine-block graph. The aim of this section is to seek ways to characterise the set of all *normalised* offset vectors (c, d) of the DAE. This also characterises the set of *valid* (c, d) , since the latter come from adding an arbitrary non-negative constant to the former.

This is not a purely theoretical exercise, since there exists at least one numerical method whose success depends on judiciously choosing non-canonical offsets instead of canonical ones.

“BTF” and “block” refer to the irreducible upper fine-block-triangular form unless said otherwise. Thus “reducible” means having more than one fine block.

5.1. Local offsets. Assume the DAE has been permuted to BTF as in the previous paragraph. *Local offsets* of part l ($l = 1, \dots, p$) are offsets obtained by structural analysis of Σ_{ll} , i.e., by treating part l as a free-standing DAE, with known driving terms coming from any coupling to other parts, which because of the BTF must have index in the range $(l+1):p$. We write a set of global offsets—*arbitrary* valid offsets of Σ as a whole—as c_1, \dots, c_n and d_1, \dots, d_n in the order of the permuted matrices in (3.1); and the *canonical* local offsets as $\widehat{c}_1, \dots, \widehat{c}_n$ and $\widehat{d}_1, \dots, \widehat{d}_n$ in the same way. (E.g., d_1, \dots, d_{N_1} and $\widehat{d}_1, \dots, \widehat{d}_{N_1}$ are the global and the canonical local offsets of the variables of part 1, and so on.)

THEOREM 5.1. *Consider an irreducible BTF of the Jacobian sparsity pattern $S_0(c, d)$ for any valid (global) offsets c, d . Then the difference between global and local offsets is constant on each part, i.e. fine block. That is, there are non-negative integers K_1, \dots, K_p such that*

$$c_i = \widehat{c}_i + K_l, \quad d_i = \widehat{d}_i + K_l \quad \text{for } i \in B_l, \text{ see (3.5).}$$

K_l is called the *lead-time of block (or part) l* .

Proof. By the definition of irreducible BTF, Σ_{ll} is irreducible as a signature matrix in its own right. By definition the local offsets, restricted to block l , are the canonical offsets of Σ_{ll} . By Lemma 4.2 (ii) the global offsets, restricted to block l , are valid offsets of Σ_{ll} . By Theorem 4.7 (ii) they differ by a block-dependent constant ≥ 0 from the local offsets. This constant is K_l . \square

5.2. The block inequalities. Assume Σ is in some irreducible fine-block form (see below (4.6)) with the blocks numbered $l = 1:p$. This may or may not put Σ , or some Jacobian sparsity pattern $S_0(c, d)$, into BTF. Our aim is to convert the basic inequalities $d_j - c_i \geq \sigma_{ij}$ between n unknowns to an equivalent set between p unknowns that are identified with the p blocks. In many cases $p \ll n$ so this makes it considerably easier to study the set of solutions.

For any general offset vector (c, d) (see (2.5)), either of c or d uniquely determines the other by the equation $d_j - c_i = \sigma_{ij}$ on a HVT, and this is independent of the HVT chosen. In what follows it is convenient to speak of “an offset vector c ”, with the other half d assumed to be defined in this way. By Theorem 5.1 the difference between c and the (unique) canonical *local* offset vector $\widehat{c} = (\widehat{c}_1, \dots, \widehat{c}_n)$ is constant on each

block. That is, there exists a vector $K = (K_1, \dots, K_p)$, depending only on c , such that $c_i - \widehat{c}_i = K_l$ whenever i is in the l th block B_l , $l = 1:p$, see (3.5). Equivalently

$$(5.1) \quad c_i = K_l + \widehat{c}_i \quad \text{for } i \in B_l, l = 1:p.$$

and correspondingly

$$(5.2) \quad d_j = K_l + \widehat{d}_j \quad \text{for } j \in B_l, l = 1:p.$$

We call K_l the *lead time*⁵ of the l th block for this vector c (and this DAE), and K the *lead time vector* of c .

Since the canonical local offsets by definition are normalised within each block, we can choose an “anchor” index $i_l \in B_l$ for each l , such that $\widehat{c}_{i_l} = 0$. Then (5.1) gives $K_l = c_{i_l}$, so that K can also be regarded as the sub-vector of c defined by

$$(5.3) \quad K = (c_{i_1}, c_{i_2}, \dots, c_{i_p}).$$

Now let σ_{ij} be an arbitrary finite entry of Σ off the block diagonal, i.e., i and j belong to different blocks, say blocks k and l . Putting (5.1, 5.2) into (2.2) gives

$$(K_l + \widehat{d}_j) - (K_k + \widehat{c}_i) \geq \sigma_{ij}.$$

Rearranging, we have reduced (2.5) to a set of inequalities in the unknowns K_l :

$$(5.4) \quad K_l - K_k \geq \sigma_{ij} - \widehat{d}_j + \widehat{c}_i.$$

Here $i \in B_k$, $j \in B_l$, and (i, j) is in the sparsity pattern S of Σ , see (4.1). That is, $(i, j) \in S_{kl}$ in the notation of (4.4).

Of all the instances of (5.4) that may exist for a given (k, l) , all are redundant except the one with the largest value on the right side. Hence define, for $k, l \in 1:p$,

$$(5.5) \quad W_{kl} = \max\{\sigma_{ij} - \widehat{d}_j + \widehat{c}_i \mid (i, j) \in S_{kl}\}.$$

Then the total set of inequalities (5.4) has been reduced to

$$(5.6) \quad K_l - K_k \geq W_{kl}, \quad \text{equivalently} \quad K_l \geq K_k + W_{kl}$$

for $k, l \in 1:p$, forming the *block inequalities* of Σ , or of the DAE.

Since every finite entry within the block-diagonal—that is, in S_{ess} —is on a HVT, all terms on the right of (5.5) are zero when $k = l$, giving the trivial relation $0 \geq 0$. If $S_{kl} = \emptyset$, it is again trivial since $W_{kl} = -\infty$. Thus in (5.6) we only need consider (k, l) for which $k \neq l$ and S_{kl} is nonempty. Equivalently, the only nontrivial contributions to (5.6) come from finite σ_{ij} outside the block diagonal, so:

The block inequalities are determined by $S \setminus S_{\text{ess}}$.

The lead times K_l are to fine blocks somewhat as offsets c_i are to individual equations⁶, so we use similar terminology:

⁵As with “offset”, the name comes from a feature of the method for solving a DAE in Taylor series.

⁶One can recast (2.2) to look like (5.6). First permute the columns of Σ to put a HVT on the main diagonal. Then a valid offset has $d_j = c_j + \sigma_{jj}$. Using this convert (2.2) to $c_j - c_i \geq w_{ij}$ where $w_{ij} = \sigma_{ij} - \sigma_{jj}$.

A solution K of (5.6) is *valid* if all $K_l \geq 0$, *normalised* if $\min_l K_l = 0$, and *canonical* if it is the elementwise smallest valid solution.

That the canonical solution exists is proved in the same way as that the canonical offsets exist.

THEOREM 5.2. *Let $\psi : \mathbb{Z}^n \rightarrow \mathbb{Z}^p$ be the map $c \mapsto K$ defined by (5.3). Then*

(i) *ψ is a bijection of the set \mathcal{C} of general offset vectors c onto the set \mathcal{K} of solutions K of (5.6).*

(ii) *Such a c is valid iff $K = \psi(c)$ is valid.*

(iii) *Such a c is normalised iff $K = \psi(c)$ is normalised.*

Proof. (i) The above construction of the W_{kl} shows that if c is a solution of (2.5) then K is a solution of (5.6), so ψ maps \mathcal{C} into \mathcal{K} . Conversely for any solution K of (5.6) construct a vector c by (5.1); then working backward we see c is a solution of (2.5) and $\psi(c) = K$, showing ψ maps \mathcal{C} onto \mathcal{K} . Since every element c_i of c is of the form (5.1) for some l we see c is uniquely determined by K , so ψ is one-to-one.

For (ii), if $c \geq 0$ then K , being a subvector of c , is also ≥ 0 . Conversely, if $K \geq 0$ then $c \geq 0$ comes from (5.1) and the fact that the canonical local offsets \widehat{c}_i are ≥ 0 .

For (iii), suppose $K = \psi(c)$ is normalised, which means it is valid and one of its components is zero. Since K is a subvector of c , and c is valid by (ii), this shows c is normalised.

Conversely suppose $c \in \mathcal{C}$ is normalised, so $c_i = 0$ for some i . Let $K = \psi(c)$, and let $l = \text{blockOf}(i)$, see (3.6), so $0 = c_i = K_l + \widehat{c}_i$ by (5.1). Then K_l and \widehat{c}_i are both ≥ 0 , by part (ii) and the definition of canonical offset. They sum to zero so must both be zero. In particular $K_l = 0$, so K is normalised. \square

We continue with the assumptions and notation of the start of this subsection. Note that in the definition of the inequalities (5.6) we considered the sparsity pattern S of Σ . Now we consider them in connection with a Jacobian sparsity pattern.

THEOREM 5.3.

(i) *Let (c, d) be a valid offset vector, with associated Jacobian sparsity pattern $S_0 = S_0(c, d)$ and lead time vector $K = \psi(c)$. Then equality holds in the k, l relation (5.6) iff the corresponding k, l sub-pattern, see (4.4), of S_0 is nonempty:*

$$K_l - K_k = W_{kl} \iff (S_0)_{kl} \neq \emptyset.$$

(ii) *$S_0(c, d)$ is in irreducible fine (upper) BTF for the given ordering of the blocks iff $K_l - K_k > W_{kl}$ whenever $k > l$.*

Proof. (i) Following the working leading to Theorem 5.2, but assuming equality, we have

$$\begin{aligned} K_l - K_k &= W_{kl} \\ \iff K_l - K_k &= \sigma_{ij} + \widehat{c}_i - \widehat{d}_j && \text{for some } i \in B_k, j \in B_l \text{ because } \geq \text{ holds} \\ &&& \text{for all such } i, j \\ \iff (K_l + \widehat{d}_j) - (K_k + \widehat{c}_i) &= \sigma_{ij} \\ \iff d_j - c_i = \sigma_{ij} = \sigma_{ij} - \sigma_{jj} && \text{using (5.1) twice} \\ \iff d_j - c_i = \sigma_{ij} && \text{using } d_j = c_j + \sigma_{jj}, \text{ as HVT is on diagonal} \\ \iff (i, j) \in S_0 && \text{for some } i \in B_k, j \in B_l, \text{ from above} \\ \iff \emptyset \neq S_0 \cap (B_k \times B_l) && \text{which } = (S_0)_{kl} \text{ by definition,} \end{aligned}$$

as required.

(ii) follows at once from (i). \square

$$\begin{array}{c}
 \Sigma = \\
 \begin{array}{c}
 f_1 \\
 f_2 \\
 f_3 \\
 f_4 \\
 f_5 \\
 f_6
 \end{array}
 \begin{array}{c}
 A \\
 B \\
 C \\
 D \\
 E \\
 F
 \end{array}
 \begin{array}{c}
 x_1 \ x_2 \ x_3 \ x_4 \ x_5 \ x_6 \\
 x \ y \ \lambda \ u \ v \ \mu \ c_i \\
 \left[\begin{array}{cccccc|c}
 2^\bullet & & 0 & & & & 4 \\
 1 & 2 & 0^\bullet & & & & 4 \\
 0 & 0^\bullet & & & & & 6 \\
 & & & 2 & & 0^\bullet & 0 \\
 & & & & 3^\bullet & 0 & 0 \\
 & & 2 & 0^\bullet & 0 & & 2
 \end{array} \right]
 \end{array}
 , \\
 \begin{array}{c}
 d_j \\
 6 \ 6 \ 4 \ 2 \ 3 \ 0
 \end{array}
 \end{array}$$

$$\begin{array}{c}
 \tilde{\Sigma} = \\
 \begin{array}{c}
 f_5 \\
 f_4 \\
 f_6 \\
 f_1 \\
 f_3 \\
 f_2
 \end{array}
 \begin{array}{c}
 E \\
 D \\
 F \\
 A \\
 C \\
 B
 \end{array}
 \begin{array}{c}
 x_5 \ x_6 \ x_4 \ x_1 \ x_2 \ x_3 \\
 v \ \mu \ u \ x \ y \ \lambda \ c_i \ \hat{c}_i \\
 \left[\begin{array}{ccc|cc|c|c}
 3 & 0 & & & & & 0 & 0 \\
 & 0 & 2 & & & & 0 & 0 \\
 0 & & 0 & & & 2 & 2 & 0 \\
 & & & 2 & 0 & & 4 & 0 \\
 & & & 0 & 0 & & 6 & 2 \\
 & & & 1 & 2 & 0 & 4 & 0
 \end{array} \right]
 \end{array}
 , \\
 \begin{array}{c}
 d_j \\
 \hat{d}_j
 \end{array}
 \begin{array}{c}
 3 \ 0 \ 2 \ 6 \ 6 \ 4 \\
 3 \ 0 \ 0 \ 2 \ 2 \ 0
 \end{array}
 \end{array}$$

FIG. 5.1. Signature matrix for Example 5.1. Above: original. Below: block form.

5.3. The fine-block graph (FBG). A convenient representation of the block inequalities (5.6) is the *fine-block graph*, FBG for short, of Σ or of the DAE. Its vertices are the irreducible blocks, enumerated $1:p$ in some arbitrary order. Each inequality (5.6) defines an edge from k to l , annotated with the weight W_{kl} . (A digraph with weights on the edges is often called a *network*.)

Thus the K values are attached to the vertices, and in going along an edge the K value must increase by at least the weight on that edge.


Example 5.1. Consider this DAE of differentiation index 7 from the companion paper [13]:

$$\begin{array}{ll}
 0 = A = x'' + x\lambda & 0 = D = u'' + u\mu \\
 (5.7) \quad 0 = B = y'' + y\lambda + (x')^3 - G & 0 = E = (v''')^2 + v\mu - G \\
 0 = C = x^2 + y^2 - L^2 & 0 = F = u^2 + v^2 - (L + c\lambda)^2 + \lambda''
 \end{array}$$

The state variables are x, y, λ, u, v , and μ ; G, L and c are constants. This artificial modification of a “two-pendula” problem was used in [12].

The signature matrix is shown in Figure 5.1, in the original row and column order and after permuting to irreducible blocks with a HVT on the main diagonal.

As noted below (5.6), to derive the block inequalities we only need consider $\tilde{\Sigma}$ entries outside the block diagonal, of which there are four. Applying (5.5) gives the following inequalities. In this example each “max” in (5.5) is over a set of only one

FIG. 5.2. *Fine-block graph of (5.7).*

element, meaning that there are no redundant inequalities to remove.

$$(5.8) \quad \begin{aligned} K_1 - K_3 &\geq -3, \\ K_2 - K_1 &\geq 0, \\ K_3 - K_2 &\geq 2, \\ K_4 - K_3 &\geq 2. \end{aligned}$$

Figure 5.2 shows the equivalent FBG. Each vertex is marked by its block index and, for convenience, the names of the equations of the block. \square

The inequalities (5.6) and resulting fine-block graph are similar to a *Critical Path Method* (CPM) diagram⁷ used in project management. In this: nodes (vertices) l are project milestones; K_l is the time at which the milestone is reached; an edge is a project activity; its weight is the minimum time needed to complete the activity; edges into a node are activities that must be completed for that milestone. One aims to minimise the time between project start and project end subject to these constraints, and one also wishes to find which activities have start-times critical for timely completion, and which have some *float* (freedom in when they are done).

Loops (cycles), and edges of negative weight, are common in a FBG. They are less so in CPM diagrams, but can represent deadlines that must be met. For instance an edge looping back from node “end” to node “start” with weight -7 says $K_{\text{start}} - K_{\text{end}} \geq -7$. So $K_{\text{end}} \leq K_{\text{start}} + 7$ —the project must take at most 7 time units.

In the context of Taylor series solution of a DAE, the FBG has a CPM-like practical meaning. E.g., a weight of 2 going from block A to block B means that for any k , the $(k+2)$ -th Taylor coefficient of variables in A must have been found before one can find the k -th coefficient of variables in B. They are thus key data for scheduling the solution process.

Example 5.2. Example 5.1 illustrates the comment after Corollary 4.6, that the ordering of blocks may be partly dictated by a specific choice of offsets c and d —equivalently, a choice of solution to the inequalities (5.6).

Writing $c = (c_A, c_B, c_C, c_D, c_E, c_F)$, i.e. with the original order of equations, and $K = (K_1, K_2, K_3, K_4)$, the map (5.3) from c to K is

$$(5.9) \quad \psi : c \mapsto K = (c_E, c_D, c_F, c_A).$$

⁷This uses the “activity on arrow” convention used by the classic PERT (Project Evaluation and Review Technique) method. Current CPM tools mostly use an “activity on node” convention.

	type X	type Y	type Z
K	$(0, 0, 2, 4 + a)$	$(0, 0, 3, 5 + a)$	$(0, 1, 3, 5 + a)$
c	$(4 + a, 4 + a, 6 + a, 0, 0, 2)$	$(5 + a, 5 + a, 7 + a, 0, 0, 3)$	$(5 + a, 5 + a, 7 + a, 1, 0, 3)$
d	$(6 + a, 6 + a, 4 + a, 2, 3, 0)$	$(7 + a, 7 + a, 5 + a, 3, 3, 0)$	$(7 + a, 7 + a, 5 + a, 3, 3, 1)$

	type X							type Y						
	v	μ	u	x	y	λ	c_i	v	μ	u	x	y	λ	c_i
E	×	×					0	×	×					0
D		×	×				0		×	×				0
F			×			●	2	×		×			●	3
A				×	×		4				×	×		5
C				×	×		6				×	×		7
B				×	×		4				×	×		5
d_j	3	0	2	6	6	4		3	0	3	7	7	5	

	type Z						
	v	μ	u	x	y	λ	c_i
E	×						0
D		×	×				1
F	×		×			●	3
A				×	×		5
C				×	×		7
B				×	×		5
d_j	3	1	3	7	7	5	

FIG. 5.3. Above: general formula for vector K and offsets, for Example 5.1, in original equation- and variable-order.

Below: corresponding Jacobian sparsity patterns, in order used for $\tilde{\Sigma}$ in Figure 5.1. In each case the entry \bullet is present if $a = 0$, absent if $a > 0$.

and its inverse that recovers c from K is

$$(5.10) \quad \psi^{-1} : K \mapsto c = (K_4, K_4, K_4 + 2, K_2, K_1, K_3).$$

By inspection one finds three normalised solutions (“types X, Y and Z”) of the block inequalities (5.8), each parameterised by an arbitrary integer $a \geq 0$, as in Figure 5.3. Every normalised solution is of this form. They are shown with the corresponding c_i (using (5.10)) and d_j (using $d_j = c_i + \sigma_{ij}$ on a HVT), in the original order of equations and variables. The corresponding Jacobian patterns $S_0(c, d)$ are also shown in Figure 5.3. Six slightly different patterns occur. Type X with $a = 0$ gives the canonical offsets, and the ordering was chosen to put this pattern into BTF. Type Y is not in BTF but can be made so by putting blocks 1, 2, 3 in the order 3, 1, 2. Similarly type Z is not in BTF but can be made so by them in the order 2, 3, 1.

Also, if $a = 0$ the entry \bullet is present forcing block 4 after block 3 in any BTF. If $a > 0$ this entry is absent, and block 4 can be in any position among the other blocks.

Some block orders are incompatible with BTF: none of the 6 types is in BTF under the order 3, 2, 1, 4. \square

Following terminology used in CPM, if K is a particular solution vector, call an edge of the FBG *critical* for K , or a *K-critical edge*, if the corresponding inequality (5.6) is satisfied as an equality. A path in the graph is a *critical path* for K , or a *K-critical path*, if all its edges are critical for K .

THEOREM 5.4. *For any solution vector K , the subgraph $\text{FBG}(K)$ formed by the K -critical edges of the FBG is acyclic.*

Proof. Let (c, d) be the offset vector such that $\psi(c) = K$, and S_0 its Jacobian sparsity pattern. S_0 can be put into irreducible BTF. By Theorem 4.5, its blocks are the (irreducible) blocks of S_{ess} in some order. Otherwise said, there exists some ordering of the blocks of S_{ess} that puts S_0 into BTF. We number the vertices of the FBG in this order.

By Theorem 5.3, for every K -critical edge $k \rightarrow l$ in the FBG, the k, l sub-pattern of S_0 is nonempty, hence by the definition of BTF it cannot be below the diagonal. That is, it must have $k < l$. Hence in this ordering, every K -critical path has a strictly decreasing sequence of vertex numbers. Thus it cannot have a cycle. \square

A linear order on the vertices of an acyclic graph *respects the graph order*, or is a *topological sort* of the graph, if u comes before v whenever there is an edge in the graph from u to v .

THEOREM 5.5. *Let (c, d) be a normalised offset vector with Jacobian sparsity pattern $S_0 = S_0(c, d)$ and lead-time vector $K = \psi(c)$. Then an ordering of the irreducible blocks puts S_0 into BTF iff it respects the order of the acyclic graph $\text{FBG}(K)$.*

Proof. This is a restatement of Theorem 5.3 (ii). \square

THEOREM 5.6. *In a FBG, the sum of the weights W_{kl} round any cycle is strictly < 0 .*

Proof. Suppose the contrary. Then there is some sequence of vertices l_0, l_1, \dots, l_m with $l_m = l_0$, where each $l_i \in 1:p$ and $m \geq 2$, such that

$$(5.11) \quad \sum_{i=1}^m W_{l_{i-1}, l_i} = 0.$$

Take an arbitrary solution K of the block-inequalities (5.6)—we know there exists at least one. Then K must satisfy them as equalities on this cycle. Namely, we have

$$(5.12) \quad K_{l_i} - K_{l_{i-1}} \geq W_{l_{i-1}, l_i}, \quad i \in 1:m.$$

If we sum all but one of these we get the reverse of the remaining one. For example, summing from the second to the last gives

$$K_{l_m} - K_{l_1} \geq \sum_{i=2}^m W_{l_{i-1}, l_i},$$

which, using (5.11), reduces to $K_{l_1} - K_{l_0} \leq W_{l_0, l_1}$, the reverse of the first of (5.12), showing $K_{l_1} - K_{l_0} = W_{l_0, l_1}$.

Hence all the edges in the cycle are K -critical and thus belong to $\text{FBG}(K)$, a contradiction since the latter is acyclic by Theorem 5.4. \square

5.4. Relation between DAE properties and FBG topology. *Strong connectedness* and *strong components* of a digraph were defined in §2.3. Our aim in this subsection is to relate connectedness properties of the FBG to properties of the DAE and of its set of normalised offset vectors. It is convenient to start with a pure graph theory result.

Let $G = (V, E)$ be a digraph and ϕ a function from V onto some set W . That is, its domain is all of V and its range is all of W . Then the sets

$$(5.13) \quad V_w = \phi^{-1}(\{w\}) = \{v \in V \mid \phi(v) = w\},$$

for $w \in W$, are a partition of V —they are nonempty, pairwise disjoint, and cover V .

DEFINITION 5.7. *The quotient graph G/ϕ has vertex set W , and its edge set F comprises all $\phi(u) \rightarrow \phi(v)$ such that $u \rightarrow v$ is an edge of G and $\phi(u) \neq \phi(v)$.*

The quotient graph can be regarded as the graph that results from collapsing each set of vertices V_w to a single point.

THEOREM 5.8. *Suppose the subgraph G_w of G generated by V_w in (5.13) is strongly connected for each w . Then*

- (i) *Vertex v is reachable from vertex u in G iff $\phi(v)$ is reachable from $\phi(u)$ in G/ϕ .*
- (ii) *The strong components of G are exactly the inverse images $\phi^{-1}(C) = \{v \in V \mid \phi(v) \in C\}$ of strong components C of G/ϕ .*

Proof. (i) Suppose there is a path in G from u to v :

$$u = v_0 \rightarrow v_1 \rightarrow \dots \rightarrow v_q = v.$$

Define the “path”

$$\phi(u) = \phi(v_0) \rightarrow \phi(v_1) \rightarrow \dots \rightarrow \phi(v_q) = \phi(v),$$

and delete from it all “edges” whose start and end are equal. This gives a genuine path (possibly of zero length) in G/ϕ from $\phi(u)$ to $\phi(v)$.

Conversely suppose there is a path in G/ϕ from $\phi(u)$ to $\phi(v)$:

$$\phi(u) = w_0 \rightarrow w_1 \rightarrow \dots \rightarrow w_q = \phi(v).$$

By definition of the quotient, each $w_{i-1} \rightarrow w_i$ has the form $\phi(v_{i-1}^+) \rightarrow \phi(v_i^-)$ where $v_{i-1}^+ \rightarrow v_i^-$ is an edge of G . For $i = 1 : q-1$, vertex v_i^- need not equal v_i^+ but they both belong to V_{w_i} , which by hypothesis is strongly connected, as a subgraph of G . Hence we can find a path π_i (possibly of zero length) from v_i^- to v_i^+ in V_{w_i} and hence in G . Glue these together to make a path

$$u = v_0^+ \rightarrow v_1^- \xrightarrow{\pi_1} v_1^+ \rightarrow v_2^- \xrightarrow{\pi_2} v_2^+ \dots \rightarrow v_{q-1}^- \xrightarrow{\pi_{q-1}} v_{q-1}^+ \rightarrow v_q^- = v$$

in G from u to v , as required.

(ii) In any digraph the relation “ v is mutually reachable from u ” (i.e., each is reachable from the other) is clearly an equivalence relation on the vertices, and the strong components are its equivalence classes.

Let C_0 be any strong component of G , let u_0 be some fixed element of C_0 , let $w_0 = \phi(u_0)$, and let C be the strong component of G/ϕ containing w_0 . For any vertex u of G

$$\begin{aligned} u \in \phi^{-1}(C) &\iff \phi(u) \in C \\ &\iff \phi(u) \text{ is mutually reachable from } w_0 = \phi(u_0) \\ &\iff u \text{ is mutually reachable from } u_0, \text{ by part (i)} \\ &\iff u \in C_0. \end{aligned}$$

Hence C_0 is the inverse image of strong component C of G/ϕ . Since C_0 was arbitrary this proves the result. \square

Now let S be the sparsity pattern of Σ and let $G = \text{Graph}(S, T)$ be the graph of S relative to some chosen HVT T (see §4.1, §4.2). Let ϕ be the map that takes each vertex v of G to the fine block it belongs to. (To be precise, regard v as a row label,

column label pair (r, c) as matched by T . It maps to l which labels the unique fine block (R_l, C_l) such that $r \in R_l, c \in C_l$.)

Then the FBG—ignoring the edge weights—is the quotient graph G/ϕ , because it has an edge from k to l iff $S_{kl} \neq \emptyset$, that is iff there is an edge in G from some vertex u in (R_k, C_k) to some vertex v in (R_l, C_l) .

The coarse blocks, by definition, are the strong components of G . By the definition of ϕ , its sets V_w in (5.13) are exactly the fine blocks, so they are the strong components of $\text{Graph}(S_{\text{ess}}, T)$ and a fortiori strongly connected in G . Hence by the last theorem, the coarse blocks are the inverse images under ϕ of the strong components of the FBG. Inverse image takes a fine block, regarded as a single vertex w of the FBG, to the same fine block regarded as a set V_w of vertices of G . Hence we have proved

THEOREM 5.9. *Each coarse block is the union of the fine blocks comprising a strong component of the FBG.*

It can be seen from Example 5.1, in particular Figure 5.3, that the general normalised offset comes from combining a finite set—the K vectors $(0, 0, 2, 4)$, $(0, 0, 3, 5)$, $(0, 1, 3, 5)$ or the derived c and d —with a parameterisation by an integer a , thus creating an infinite set of offsets. The finite set is associated with a cycle in the FBG forming a strong component; the parameter is associated with an edge between two strong components. The next result goes some way toward explaining this behaviour.

THEOREM 5.10.

- (i) *The set of normalised offset vectors has more than one member iff the FBG has more than one vertex, i.e. iff the DAE is fine-block reducible.*
- (ii) *The set of normalised offset vectors is infinite iff the FBG has more than one strong component, i.e. iff the DAE is coarse-block reducible.*

Proof. (i) If the DAE is irreducible, there is only one normalised offset vector by Theorem 4.7. Conversely suppose the DAE is reducible, so the FBG has at least two vertices. There is at least one normalised offset vector (c, d) , namely the canonical one. Consider the corresponding canonical $K = \psi(c)$, and the acyclic subgraph $\text{FBG}(K)$ it defines. Being acyclic, this graph has at least one final vertex k with no edges going out of it. Then increasing K_k to $K_k + 1$, leaving the rest of K unchanged, gives a new valid $K' \neq K$. For any edge in the FBG leaving k , say $k \rightarrow l$, cannot be K -critical else it would be part of $\text{FBG}(K)$. Thus the value $K_l - K_k - W_{kl}$ is > 0 and being an integer is at least 1. Hence changing K to K' cannot violate the inequality. In case K' is not normalised, subtracting 1 from all its entries makes it so. This proves the converse.

(ii) Strongly connected means that for any two vertices k, l there is a path in the graph from k to l , and one from l to k . Summing the inequalities (5.6) along such a path shows there is an upper bound M (e.g. the sum of the absolute values of the weights on all the edges) for $K_l - K_k$ and also $K_k - K_l$, hence for $|K_l - K_k|$, independent of K, k and l . A normalised $K = (K_1, \dots, K_p)$ has integer components $K_i \geq 0$, one of which is 0, so $0 \leq K_i \leq M$. Hence there are at most M^p of them.

Conversely if the FBG is not strongly connected it breaks into more than one strong component. Hence the vertices $1:p$ can be partitioned into two nonempty subsets V_1, V_2 defining subgraphs G_1 and G_2 , such that there is no edge (in the whole FBG) from G_2 to G_1 . Then for any integer $a > 0$ one can add a to K_l for each $l \in V_2$ to get $K'(a)$ —again, normalising by subtracting a constant if necessary. It is easily seen that each resulting $K'(a)$ is a normalised solution and $K'(a) \neq K'(b)$ whenever $a \neq b$, thus giving infinitely many normalised solutions. \square

6. Conclusion.

6.1. Application of the theory. Much of the theory presented here is implemented in the current version of our structural analysis tool DAESA. It can read MATLAB code for the DAE, or it can input a signature matrix directly. It can display (all or a section of) the original signature matrix, or the result of permuting to coarse or fine block form. Among the functionality we plan for future releases is:

- A numerical check that the Jacobian is nonsingular at some chosen point. This will let a user check if the Σ -method succeeds in the sense defined in §2.2.
- A stage-wise description, on the lines of DAESA’s current “solution scheme” function, of index reduction of a DAE by the dummy derivatives method—to be described in a separate paper.
- Functions to analyse the DAE via its fine-block graph on the lines of Section 5.

6.2. Computational costs. Permuting a known sparsity pattern to BTF, typically by an implementation of the Dulmage–Mendelsohn decomposition, is quite cheap. Pothen and Fan [15] give a complexity analysis and comparison of several methods.

Computing Σ , hence its sparsity pattern S , is straightforward from suitable program code describing the DAE, but see the caveats in §4.3.

To find a Jacobian sparsity pattern one needs to have solved the LAP first to find offsets c, d . Once that is done, computing $S_0(c, d)$ is simple, given Σ .

Solving the LAP itself appears the most challenging aspect, especially for large problems. The best current methods have complexity $O(n^3)$ for dense problems and $O(n^2 \log n + nm)$ using sparsity, where m is the number of nontrivial entries (finite σ_{ij}), see [2].

It looks desirable to find the essential sparsity pattern S_{ess} instead, since this, like S , is an inherent property of the DAE independent of any offsets. However we do not know of any better method than to find some offsets c, d first and the BTF of $S_0(c, d)$, and then to use Theorem 4.5.

Hence our current approach is a “bootstrap” one. We find the coarse BTF based on S . On each coarse block, we solve the LAP and construct a HVT and “coarse” local offsets. From these we can efficiently compute, independently of each other, (a) global offsets for the whole problem and (b) fine blocks within each coarse block and their “fine” local offsets. This gives scope for parallelism on large problems.

The success of this method depends heavily on problem structure, in particular the size distribution of coarse blocks. In future work we aim to present systematic results about how it performs on large problems.

REFERENCES

- [1] R. BALAKRISHNAN AND K. RANGANATHAN, *A textbook of graph theory*, Springer, 2012.
- [2] R. E. BURKARD, M. DELL’AMICO, AND S. MARTELLO, *Assignment Problems, Revised Reprint*, SIAM, 2009.
- [3] S. L. CAMPBELL AND E. GRIEPENTROG, *Solvability of general differential algebraic equations*, SIAM Journal on Scientific Computing, 16 (1995), pp. 257–270.
- [4] T. F. COLEMAN, A. EDENBRANDT, AND J. R. GILBERT, *Predicting fill for sparse orthogonal factorization*, J. ACM, 33 (1986), pp. 517–532.
- [5] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct Methods for Sparse Matrices*, Oxford Science Publications, Clarendon Press, Oxford, 1986.
- [6] A. GRIEWANK AND A. WALTHER, *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*, SIAM, Philadelphia, PA, USA, second ed., 2008.

- [7] B. L. *Mathematical Programming in Practice*, Topics in Operational Research, Sir Isaac Pitman & Sons Limited, 1968.
- [8] S. E. MATTSSON AND G. SÖDERLIND, *Index reduction in differential-algebraic equations using dummy derivatives*, SIAM J. Sci. Comput., 14 (1993), pp. 677–692.
- [9] K. MURTY, *Linear programming*, Wiley, 1983.
- [10] N. S. NEDIALKOV AND J. D. PRYCE, *Solving differential-algebraic equations by Taylor series (I): Computing Taylor coefficients*, BIT, 45 (2005), pp. 561–591.
- [11] N. S. NEDIALKOV AND J. D. PRYCE, *Solving differential algebraic equations by Taylor series (III): the DAETS code*, JNAIAM J. Numer. Anal. Indust. Appl. Math, 3 (2008), pp. 61–80.
- [12] N. S. NEDIALKOV, J. D. PRYCE, AND G. TAN, *DAESA — a Matlab tool for structural analysis of DAEs: Software*, ACM Transactions on Mathematical Software, to appear (2014). 15 pages.
- [13] N. S. NEDIALKOV, G. TAN, AND J. D. PRYCE, *Exploiting fine block triangularization and quasi-linearity in differential-algebraic equation systems*. Submitted to SIAM J. Sci. Comput., November 2014, 18 pages.
- [14] C. C. PANTELIDES, *The consistent initialization of differential-algebraic systems*, SIAM. J. Sci. Stat. Comput., 9 (1988), pp. 213–231.
- [15] A. POTHEN AND C.-J. FAN, *Computing the block triangular form of a sparse matrix*, ACM Transactions on Mathematical Software, 16 (1990), pp. 303–324.
- [16] J. D. PRYCE, *A simple structural analysis method for DAEs*, BIT, 41 (2001), pp. 364–394.
- [17] J. D. PRYCE, N. S. NEDIALKOV, AND G. TAN, *DAESA — a Matlab tool for structural analysis of DAEs: Theory*, ACM Transactions on Mathematical Software, to appear (2014). 20 pages.
- [18] A. A. SHALENINOV, *Removal of topological degeneracy in systems of differential evolution equations*, Automation and Remote Control, 51 (1991), pp. 1599–1605 (English version). Translation from *Avtomatika i Telemekhanika*.