On the Block Triangular Form of Symmetric Matrices

Iain S. Duff, Bora Uçar

ON THE BLOCK TRIANGULAR FORM OF SYMMETRIC MATRICES*

IAIN S. DUFF[†‡] AND BORA UÇAR[§]

Abstract. We present some observations on the block triangular form (btf) of structurally symmetric, square, sparse matrices. If the matrix is structurally rank deficient, its canonical btf has at least one underdetermined and one overdetermined block. We prove that these blocks are transposes of each other. We further prove that the square block of the canonical btf, if present, has a special fine structure. These findings help us recover symmetry around the anti-diagonal in the block triangular matrix. The uncovered symmetry helps us to permute the matrix in a special form which is symmetric along the main diagonal while exhibiting the blocks of the original btf. As the square block of the canonical btf has full structural rank, the observation relating to the square block applies to structurally nonsingular, square symmetric matrices as well.

Key words. sparse matrices, block triangular form, Dulmage-Mendelsohn decomposition, maximum cardinality matchings

AMS subject classifications. 05C50, 05C70, 05D15, 65F50

1. Introduction.

We are interested in the block triangular form (btf) of structurally symmetric sparse matrices. We include in our study the particular case when a matrix is structurally rank deficient. Throughout the paper, A is always a structurally symmetric matrix with no all-zero rows or columns. The block triangular form is based on a canonical decomposition of bipartite graphs known as the Dulmage-Mendelsohn decomposition [7] (see [14] for a detailed account). When permuted into the block triangular form, the matrix A assumes the form

$$\begin{array}{c} \\ H_R \\ S_R \\ V_R \end{array} \begin{array}{c} \begin{array}{ccc} H_C & S_C & V_C \end{array} \\ \begin{pmatrix} A_H & * & * \\ O & A_S & * \\ O & O & A_V \end{pmatrix} \end{array}. \tag{1.1}$$

As we shall see, the three blocks on the diagonal are of special importance. The block A_H, formed by the rows in the set H_R and the columns in the set H_C, is underdetermined; the block A_S, formed by the rows in the set S_R and the columns in the set S_C, is square and can itself have a btf with the blocks on the diagonal being all square; the block A_V, formed by the rows in the set V_R and the columns in the set V_C, is overdetermined. As in [14], we will call these three blocks horizontal, square, and vertical, respectively. In (1.1), there are no nonzero entries in the sub-diagonal blocks shown as O. If A is structurally rank deficient, the horizontal and vertical blocks are non-empty, otherwise those blocks are empty.

Consider the following 3×3 symmetric, structurally rank deficient matrix and its

*This work was supported by "Agence Nationale de la Recherche", through the SOLSTICE project ANR-06-CIS6-010.

[†]Atlas Centre, RAL, Oxon, OX11 0QX, England (i.s.duff@rl.ac.uk).

[‡]CERFACS, 42 Av. G. Coriolis, 31057, Toulouse, France (duff@cerfacs.fr)

[§]Centre National de la Recherche Scientifique, Laboratoire de l'Informatique du Parallélisme, (UMR CNRS-ENS Lyon-INRIA-UCBL), Université de Lyon, 46, allée d'Italie, ENS Lyon, F-69364, Lyon, France (bucar@ens-lyon.fr).

btf shown on the right (we provide a summary on how to obtain the btf later):

$$\begin{array}{c} \\ r_1 \\ r_2 \\ r_3 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} c_1 & c_2 & c_3 \end{array} \\ \begin{pmatrix} \times & \times & \times \\ \times & 0 & 0 \\ \times & 0 & 0 \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ r_1 \\ r_2 \\ r_3 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} c_3 & c_2 & c_1 \end{array} \\ \begin{pmatrix} \times & \times & \times \\ 0 & 0 & \times \\ 0 & 0 & \times \end{pmatrix} \end{array}. \tag{1.2}$$

In this matrix, the square block is empty; the horizontal block has a single row r_1 and two columns c_2 and c_3; the vertical block has a single column c_1 and two rows r_2 and r_3. Notice that the set of row indices in H_R is equal to the set of column indices in V_C, furthermore the set of column indices in H_C is equal to the set of row indices in V_R. In other words, the blocks A_H and A_V for this matrix are structural transposes of each other. One of the main results in this paper is that this last relation between the horizontal and vertical blocks of the btf holds for any structurally rank deficient, symmetric matrix.

Consider now the following 3×3 symmetric, structurally full rank matrix with its btf shown on the right

$$\begin{array}{c} \\ r_1 \\ r_2 \\ r_3 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} c_1 & c_2 & c_3 \end{array} \\ \begin{pmatrix} \times & \times & \times \\ \times & \times & 0 \\ \times & 0 & 0 \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ r_1 \\ r_2 \\ r_3 \end{array}\!\!\begin{array}{c} \begin{array}{ccc} c_3 & c_2 & c_1 \end{array} \\ \begin{pmatrix} \times & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \end{pmatrix} \end{array}. \tag{1.3}$$

As the matrix is full rank, the btf has only a square block which in turn has three square blocks on the diagonal. The first block has row r_1 and column c_3; the second block has row r_2 and column c_2; the third block has row r_3 and column c_1. In this example, a block has the set of row indices equal to the set of column indices (as in the second block in the example), or the sets of row indices and column indices are disjoint and there is another block whose sets of row indices and column indices are respectively equal to the sets of column indices and row indices of the former block (as in the first and third blocks in the example). One of the main results in this paper is that this last observation on the square blocks of the btf holds for any symmetric, structurally full rank matrix.

In the following two subsections, we provide the reader with definitions (mostly standard), and necessary background material from Duff, Erisman, and Reid [4, Chapter 6], Pothen and Fan [14], and Pothen [13, Section 2.7] on the computation and properties of the btf. The computation of the btf is based on maximum cardinality matchings, or just maximum matchings, in bipartite graphs (these are discussed in Sections 1.1 and 1.2). We discuss two transformations on maximum matchings of symmetric matrices in Section 2. One of the transformations is used in [5]; the other is based on the notion of cycles of a permutation, and to the best of our knowledge is discussed and used for the first time in this paper. We use these transformations to show that for a symmetric matrix there is a maximum matching with some special properties. In Section 3, we formally state the main results on the btf of symmetric matrices that we illustrated in the examples (1.2) and (1.3),

1.1. Definitions. As is common, we associate a bipartite graph $G = (R \cup C, E)$ with the $n \times n$ matrix A, where $R = \{r_1, \ldots, r_n\}$ and $C = \{c_1, \ldots, c_n\}$ are the two sets of the vertex bipartition, and E is the set of edges. Here, the vertices in R and C correspond to the rows and the columns of A, respectively, such that $(r_i, c_j) \in E$ if and only if $a_{ij} \neq 0$. For a given $i \in \{1, \ldots, n\}$, the row r_i and the column c_i

are referred to as symmetric counterparts of each other. Similarly, the edges (r_i, c_j) and (r_j, c_i) are called symmetric counterparts of each other. When necessary, we will make it clear whether a vertex is a row or a column vertex, usually by labelling the vertex with an r or a c respectively. An edge $(r_i, c_j) \in E$ is said to be incident on the vertices r_i and c_j. Two vertices are called adjacent if there is an edge incident on both. The set of vertices that are adjacent to a vertex v is indicated by $\text{adj}(v)$. A path is a sequence of edges of the form $\big((v_0, v_1), (v_1, v_2), \ldots, (v_{k-1}, v_k)\big)$. A cycle is a sequence of edges of the form $\big((v_0, v_1), (v_1, v_2), \ldots, (v_{k-1}, v_k)\big)$ where $v_k = v_0$.

A set of edges $\mathcal{M}$ is a matching if no two edges in $\mathcal{M}$ are incident on the same vertex. In matrix terms, a matching corresponds to a set of nonzero entries no two in the same row or column. A vertex is said to be matched (with respect to a given matching) if there is an edge in the matching incident on the vertex, and to be unmatched otherwise.

Given a matching $\mathcal{M}$, an $\mathcal{M}$-alternating path is a path whose edges are alternately in $\mathcal{M}$ and not in $\mathcal{M}$. We use the notation $u \xrightarrow{\mathcal{M}} v$ to denote that vertex u reaches vertex v with an $\mathcal{M}$-alternating path (we assume that u and v are different, in other words an alternating path has at least one edge). Note that this is a bidirectional relation in an undirected graph: if $u \xrightarrow{\mathcal{M}} v$, then $v \xrightarrow{\mathcal{M}} u$. An alternating path is called an augmenting path, if it starts and ends at unmatched vertices.

The cardinality of a matching is the number of edges in it. A maximum cardinality matching or a maximum matching is a matching of maximum cardinality. Given a bipartite graph G and a matching $\mathcal{M}$, a necessary and sufficient condition for $\mathcal{M}$ to be of maximum cardinality is that there is no $\mathcal{M}$-augmenting path in G (the result is due to Berge [1] and is also summarized in different places, see for example [10, Chapter 1]).

We use $\text{mate}(v)$, to denote the vertex matched to the vertex v in a matching $\mathcal{M}$, e.g., if $\text{mate}(r_i) = c_j$, then we also have $\text{mate}(c_j) = r_i$. We use $\langle \cdot, \cdot \rangle$ to differentiate a matching edge from an ordinary edge, e.g., we use $\langle r_i, c_j \rangle$ or $\langle c_j, r_i \rangle$ to denote that the row r_i is matched to the column c_j. We say a vertex set X is completely matched to another one Y, if for all $x \in X$, we have $\text{mate}(x) \in Y$; for clarity we note that $|X| \leq |Y|$, where $|\cdot|$ denotes the cardinality of a set.

Some of the definitions in this paragraph can be found in [11]. Let A be an $n \times n$ matrix, and $\mathcal{I}$ and $\mathcal{J}$ be two subsets of $\{1, \ldots, n\}$. The matrix formed by selecting the rows and columns indexed by $\mathcal{I}$ and $\mathcal{J}$, respectively, is called a submatrix of A confined to the rows in $\mathcal{I}$ and the columns in $\mathcal{J}$. The matrix A is said to be partly decomposable if it contains an $s \times (n - s)$ zero submatrix, for $s \geq 1$. More explicitly, A is partly decomposable if there exist permutation matrices P and Q such that

$$PAQ = \begin{pmatrix} B & C \\ O & D \end{pmatrix},$$

with B and D being square. If A contains no $s \times (n - s)$ zero submatrix for $s = 1, \ldots, n - 1$, then it is called fully indecomposable, also called irreducible [4]. We note for later use that an $n \times n$ symmetric matrix A, where $n > 2$ and n is odd, $a_{ij} = a_{ji} \neq 0$ for $i = 1, \ldots, n$ and $j \equiv i + 1 \pmod{n}$, and $a_{ij} = a_{ji} = 0$ elsewhere, is fully indecomposable. The bipartite graph of this matrix is a cycle on $2n$ vertices with n row vertices and n column vertices. Any $n \times n$ matrix B whose sparsity structure is a superset of that of A, i.e., $b_{ij} \neq 0$ if $a_{ij} \neq 0$, is also fully indecomposable.

1.2. Computation and properties of the block triangular form. Algorithms to compute the btf of a matrix have been available for some time, see for

example [6, 8, 14]. The first step of these algorithms is to find a maximum cardinality matching $\mathcal{M}$ on the bipartite graph representation of the given matrix. Then, the row and column vertices are classified using the following equations based on alternating paths with respect to $\mathcal{M}$ (the equations are rephrased from [13, Section 2.7]):

$$
\begin{aligned}
U_C &= \{c \in C : c \text{ is unmatched}\} & U_R &= \{r \in R : r \text{ is unmatched}\} \\
H_R &= \{r \in R : r \xrightarrow{\mathcal{M}} u \text{ for some } u \in U_C\} & V_C &= \{c \in C : c \xrightarrow{\mathcal{M}} u \text{ for some } u \in U_R\} \\
H'_C &= \{c \in C : c \xrightarrow{\mathcal{M}} u \text{ for some } u \in U_C\} & V'_R &= \{r \in R : r \xrightarrow{\mathcal{M}} u \text{ for some } u \in U_R\} \\
H_C &= U_C \cup H'_C & V_R &= U_R \cup V'_R \\
S_R &= R \setminus (H_R \cup V_R) & S_C &= C \setminus (H_C \cup V_C)\ .
\end{aligned}
$$

For completeness, we provide Algorithm 1 that finds the set of rows H_R and the set of columns H_C of the horizontal block. The algorithm grows the row set H_R and the column set H_C by running a graph search algorithm. At a column vertex c known to be in H_C (whose adjacency is not explored yet), it adds all neighbouring rows to H_R. At a row vertex r known to be in H_R, the algorithm only visits the column $v = \text{mate}(r)$ and adds v to H_C if it is not already there.

Algorithm 1 Algorithm to find horizontal block of the btf

Input $\mathcal{M}$: a maximum matching; mate(v) giving the mate of a vertex v
Output H_R and H_C: the set of rows and columns of the horizontal block

```
1: H_R ← H_C ← ∅
2: U ← {unmatched columns}
3: while U ≠ ∅ do
4:     Pick a column vertex c ∈ U and set U ← U \ {c}
5:     H_C ← H_C ∪ {c}
6:     for each row vertex r ∈ adj(c) \ H_R do
7:         H_R ← H_R ∪ {r}
8:         v ← mate(r)      ▷ should exist; otherwise flags an augmenting path
9:         if column v ∉ U ∪ H_C then
10:            U ← U ∪ {v}
```

A similar algorithm is run to find the rows and columns in the vertical block. In this case, at a column vertex in V_C, only its mate is visited and added to V_R, if necessary; at a row vertex in V_R, the neighbouring columns are added to V_C, whenever necessary. After finding the rows and columns of the horizontal and vertical blocks, the remaining rows and columns are marked to be in the sets S_R and S_C, respectively.

We note the following properties of the block triangular form without proving them. The proofs can be found in [3], [4, Chapter 6] and [13, Section 2.7]. These properties hold for any matrix.

FACT 1.1. *The rows in H_R are completely matched to the columns in H_C. The columns in S_C are completely matched to the rows in S_R and vice versa. The columns in V_C are completely matched to the rows in V_R.*

FACT 1.2. *The block triangular form is unique. In other words, any maximum matching yields the same sets H_R, H_C, S_R, S_C, V_R, and V_C.*

The previous two properties also imply that all entries of a maximum matching should reside in the blocks on the diagonal of the btf.

FACT 1.3. *In the block triangular form of a structurally rank deficient, square*

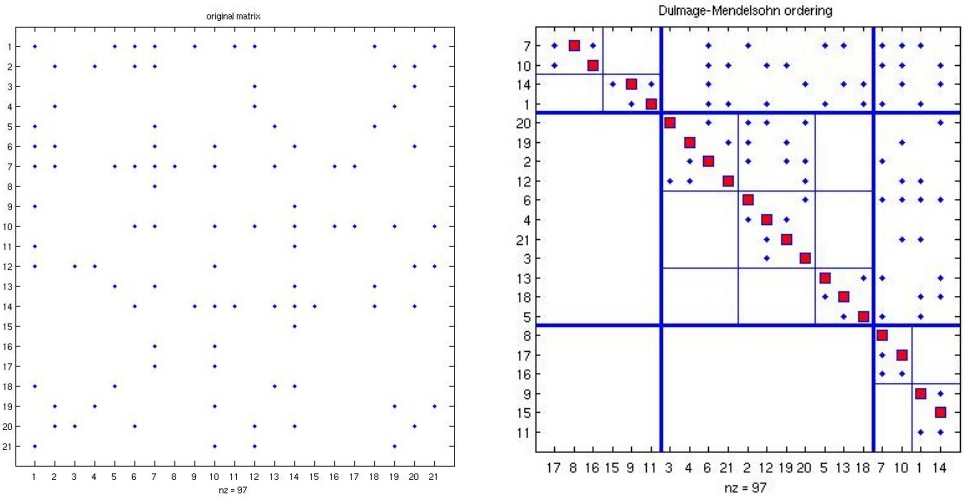

FIG. 1.1. *A symmetric* 21×21 *matrix with a structural rank of 19 and its Dulmage-Mendelsohn ordering induced by a maximum cardinality matching shown with highlighted nonzeros.*

matrix (not necessarily symmetric), the horizontal and vertical blocks both should be present. The square block may be missing.

It may be possible to decompose the three blocks on the diagonal A_H, A_S, and A_V further into smaller submatrices, resulting in a fine decomposition. In the fine decomposition, the horizontal and vertical blocks have block diagonal structure where the individual blocks on the diagonal are horizontal and vertical, respectively. In a given block, each vertex (row or column vertex) can reach all the other vertices in the bipartite graph representation of the same block.

The fine decomposition of the square block A_S is obtained by identifying irreducible blocks. Essentially, the fine decomposition is achieved by including all rows and columns corresponding to the vertices in an alternating tour (an alternating cycle where one can repeat edges) in a single block. We summarise some of the properties of the fine decomposition of the square block that will be needed in the rest of the paper. Let p be the number of irreducible blocks (all of them square) in the fine decomposition of A_S, and let R_i and C_i be the sets of row and column indices in the ith block for $i = 1, \ldots, p$. The rows whose indices are in R_i are matched to the columns whose indices are in C_i. The sets R_i and C_i for $i = 1, \ldots, p$ are unique—they are independent of the choice of maximum matching [3]. The blocks cannot be combined to yield another decomposition satisfying the properties given above.

Figure 1.1 shows a symmetric matrix and the fine decomposition induced by a maximum matching shown with highlighted nonzero entries. There are two fine blocks in the horizontal and vertical submatrices and three blocks in the square submatrix. Horizontal and vertical lines are used to draw the borders of the fine blocks, and thicker lines are used to separate the square coarse block from the others.

2. Two transformations. As discussed in the previous section, the btf is unique and can be obtained using any maximum matching. In this section, we discuss two processes that transform any maximum matching in the bipartite graph of a symmetric

matrix into another that has some special properties. We first need a few definitions.

DEFINITION 2.1. *Given a matching, not necessarily of maximum cardinality, in the bipartite graph of a square matrix,*

an m**-sequence** *is a sequence of edges* $\big(\langle r_i, c_j\rangle, \langle r_j, \cdot\rangle, \ldots, \langle \cdot, c_k\rangle, \langle r_k, c_\ell\rangle\big)$*, where each edge is in the matching. Notice that an edge* $\langle r_k, c_\ell\rangle$ *in an* m*-sequence is followed by a matching edge of the form* $\langle r_\ell, \cdot\rangle$*.*

a closed m**-sequence** *is an* m*-sequence where the start and end vertices are symmetric counterparts of each other.*

an open m**-sequence** *is an* m*-sequence which is not closed.*

Note that the above definitions depend on the numbering of the rows and the columns and hence it is important to use a numbering (in our case the original one) that preserves the symmetric structure.

We note that an m-sequence is not necessarily a path in the bipartite graph, as the two consecutive matching edges $\langle \cdot, c_k\rangle$ and $\langle r_k, \cdot\rangle$ are not necessarily connected by the edge (c_k, r_k). Without loss of generality, we assume that an m-sequence starts at a row vertex and ends at a column vertex. Consider the matching shown in Fig. 1.1. In this figure, $\big(\langle r_{14}, c_9\rangle, \langle r_9, c_1\rangle, \langle r_1, c_{11}\rangle\big)$ forms an m-sequence. Any matching in the bipartite graph of a square matrix can be decomposed into open and closed m-sequences. In order to see this, consider the following process. Take a matching edge $\langle r_i, c_j\rangle$, which is an m-sequence starting at row vertex r_i and ending at column vertex c_j; if the symmetric counterpart c_i of the starting vertex is also matched, say to r_k, extend the m-sequence by the matching edge $\langle r_k, c_i\rangle$ to obtain the m-sequence $(\langle r_k, c_i\rangle, \langle r_i, c_j\rangle)$; if the symmetric counterpart r_j of the end vertex is also matched, say to c_ℓ, extend the m-sequence by the matching edge $\langle r_j, c_\ell\rangle$ and continue the process until the symmetric counterparts of the start and end vertices are also in the path so grown (obtaining a closed m-sequence) or they are both unmatched (obtaining an open m-sequence). Consider the m-sequence given above. Since the symmetric counterpart r_{11} of the end vertex is not matched, we cannot add another edge to the end. Since, the symmetric counter part c_{14} of the starting vertex is matched, we add $\langle r_{15}, c_{14}\rangle$ to the beginning. The symmetric counterpart c_{15} of the new starting vertex is not matched, therefore we obtain an open m-sequence $OS_1 = \big(\langle r_{15}, c_{14}\rangle, \langle r_{14}, c_9\rangle, \langle r_9, c_1\rangle, \langle r_1, c_{11}\rangle\big)$. The other open m-sequence is $OS_2 = \big(\langle r_{17}, c_{10}\rangle, \langle r_{10}, c_{16}\rangle\big)$. There are five closed m-sequences: the first is $CS_1 = \big(\langle r_7, c_8\rangle, \langle r_8, c_7\rangle\big)$; the second is $CS_2 = \big(\langle r_{20}, c_3\rangle, \langle r_3, c_{20}\rangle\big)$; the third is $CS_3 = \big(\langle r_{19}, c_4\rangle, \langle r_4, c_{12}\rangle, \langle r_{12}, c_{21}\rangle, \langle r_{21}, c_{19}\rangle\big)$; the fourth is $CS_4 = \big(\langle r_2, c_6\rangle, \langle r_6, c_2\rangle\big)$; and the fifth is $CS_5 = \big(\langle r_{13}, c_5\rangle, \langle r_5, c_{18}\rangle, \langle r_{18}, c_{13}\rangle\big)$.

2.1. Automorphic maximum matchings. We define a matching to be automorphic if it matches a set of rows to the corresponding set of columns. That is, for a matching $\mathcal{M}$ to be automorphic, whenever $\langle r_i, c_j\rangle \in \mathcal{M}$, the column c_i and the row r_j should be matched by $\mathcal{M}$—not necessarily to each other.

We need the following theorem [9, Theorem 4.1, p. 191] which is a restatement of a result by Mendelsohn and Dulmage [12, Theorem 1].

THEOREM 2.2. *Let* $G = (R \cup C, E)$ *be a bipartite graph and let* $\mathcal{M}_1, \mathcal{M}_2$ *be two matchings in* G*. Then, there exists a matching* $\mathcal{M} \subseteq \mathcal{M}_1 \cup \mathcal{M}_2$*, such that* $\mathcal{M}$ *covers all the vertices of* R *covered by* $\mathcal{M}_1$ *and all the vertices of* C *covered by* $\mathcal{M}_2$*.*

Let $\mathcal{M}$ be a maximum cardinality matching in the bipartite graph of a symmetric matrix, and $\mathcal{I}$ and $\mathcal{J}$ be the set of row and column indices matched by $\mathcal{M}$, i.e., $\mathcal{I} = \{i : \langle r_i, c_j\rangle \in \mathcal{M} \text{ for some } j\}$ and $\mathcal{J} = \{j : \langle r_i, c_j\rangle \in \mathcal{M} \text{ for some } i\}$. Then, because of the symmetry of the matrix, the matching $\mathcal{M}' = \{\langle r_j, c_i\rangle : \langle r_i, c_j\rangle \in \mathcal{M}\}$ is

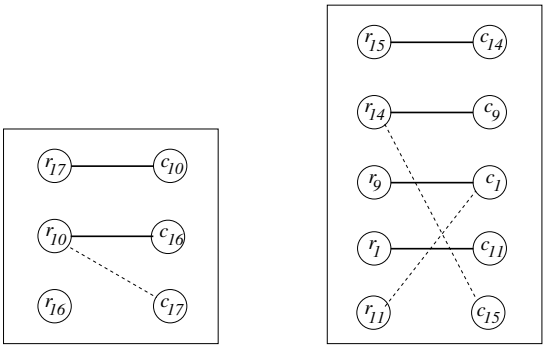

FIG. 2.1. *The two open m-sequences are shown with solid lines. The set of row indices and column indices are extended to have the same set of row and column vertices. Dashed lines are symmetric counterparts of some matching edges; they are used to update the matching such that the new matching is automorphic. The new matching edges for the left subfigure is* $\big(\langle r_{17}, c_{10}\rangle, \langle r_{10}, c_{17}\rangle\big)$, *and the new matching for the right subfigure is* $\big(\langle r_{15}, c_{14}\rangle, \langle r_{14}, c_{15}\rangle, \langle r_1, c_{11}\rangle, \langle r_{11}, c_1\rangle\big)$.

also of maximum cardinality and has the set of row indices $\mathcal{J}$ and the set of column indices $\mathcal{I}$. By the application of Theorem 2.2, there is a matching that covers the set of rows and columns indexed by $\mathcal{I}$ and is of maximum cardinality as $|\mathcal{I}|$ is the size of a maximum matching. That is, there is a maximum cardinality matching which covers the same set of row and column indices. We note that this last result is proved by Duff and Pralet [5, Property 4.2] without using Theorem 2.2.

Consider the matching given in Fig. 1.1 and the open m-sequences OS_1 and OS_2 found above. The open m-sequences are shown in Fig. 2.1, OS_1 on the left and OS_2 on the right, with solid lines along with some edges (dashed lines) that are symmetric counterparts of some matching edges. The dashed edges are used to modify the matching so that the edge $\langle r_{10}, c_{17}\rangle$ replaces $\langle r_{10}, c_{16}\rangle$ for OS_1 become a closed m-sequence, $\langle r_{14}, c_{15}\rangle$ replaces $\langle r_{14}, c_9\rangle$, and $\langle r_{11}, c_1\rangle$ replaces $\langle r_9, c_1\rangle$ for OS_2 become a closed m-sequence. Combined with the existing closed m-sequences, the new maximum matching becomes automorphic as shown in Fig. 2.2, as is easily seen by observing that neither the rows nor the columns 9 and 16 is in the matching.

2.2. Permutation cycles of an automorphic matching. An automorphic matching from $\mathcal{I}$ to $\mathcal{I}$ can be perceived as a permutation of the set $\mathcal{I}$ in an algebraic sense (a one-to-one and onto function). By starting from an element of the set $\mathcal{I}$ and by applying the permutation until the starting element is seen again, we can obtain cycles of the permutation (for more on cycles of a permutation see [2, Section 1.5]). Similarly, by following the matching edges of an automorphic matching $\mathcal{M}$ as $\big(\langle r_i, c_j\rangle, \langle r_j, \cdot\rangle, \ldots, \langle \cdot, c_i\rangle\big)$, we can obtain the closed m-sequences of $\mathcal{M}$. Due to this correspondence, we refer to the closed m-sequences of an automorphic matching as the permutation cycles.

Note that as the permutation cycles are also closed m-sequences, they do not necessarily correspond to ordinary cycles in the underlying bipartite graph. Consider for example the permutation cycle $\big(\langle r_i, c_j\rangle, \langle r_j, c_i\rangle\big)$. If a_{ii} or a_{jj} is zero in A, then we do not have a cycle in the graph; we only have a permutation cycle. The length of a permutation cycle is the number of matching edges in it. The length 2 permutation cycles, also called transpositions [11, p.11], are of the form $\big(\langle r_i, c_j\rangle, \langle r_j, c_i\rangle\big)$ and are of special importance. Figure 2.3 displays permutation cycles of length 1 to 4 in

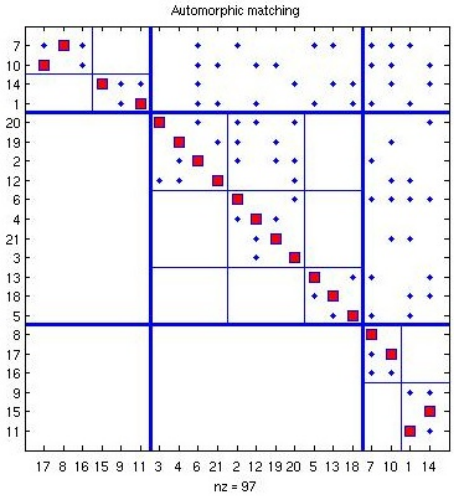

FIG. 2.2. *An automorphic maximum matching is shown with highlighted nonzeros.*

a hypothetical example. An odd permutation cycle is of odd length, and an even permutation cycle is of even length. Note that the edges of an odd permutation cycle (with length greater than one) when put together with their symmetric counterparts form a unique ordinary cycle in the bipartite graph of A. Note also that any pair of a row and a column in an odd permutation cycle is reachable from each other via two alternating paths: one starting and ending with a matching edge, the other starting and ending with an unmatched edge. For an even permutation cycle of length k, adding the symmetric edges partitions the permutation cycle into two ordinary cycles each having $k/2$ matching edges and $k/2$ non-matching edges, where the row vertices in one cycle are symmetric counterparts of the column vertices in the other one; the two cycles may be connected in the bipartite graph due to existence of other edges, but we are not interested in this possibility. Consider, for example, the length 4 permutation cycle $\big(\langle r_i, c_j\rangle,\langle r_j, c_k\rangle, \langle r_k, c_l\rangle, \langle r_l, c_i\rangle\big)$ shown in Fig. 2.3. The permutation cycle is split between two ordinary cycles in the bipartite graph of A: $\big(\langle r_i, c_j\rangle,(c_j, r_k),\langle r_k, c_l\rangle,(c_l, r_i)\big)$ and $\big(\langle r_j, c_k\rangle,(c_k, r_l), \langle r_l, c_i\rangle,(c_i, r_j)\big)$. As seen, each of these cycles contain 2 matching edges and 2 non-matching edges, and the row vertices in one cycle are symmetric counterparts of the column vertices in the other.

From an automorphic matching $\mathcal{M}$, we construct another automorphic matching $\mathcal{M}'$ which is composed of odd length permutation cycles and length 2 permutation cycles. We proceed as follows. First, all edges of $\mathcal{M}$ that form an odd permutation cycle are copied into $\mathcal{M}'$, e.g., for a length 3 permutation cycle of the form $\big(\langle r_i, c_j\rangle, \langle r_j, c_k\rangle, \langle r_k, c_i\rangle\big)$, these three edges are copied into $\mathcal{M}'$. Then, even length permutation cycles of $\mathcal{M}$ are decomposed into length 2 permutation cycles, and these length 2 permutation cycles are added to $\mathcal{M}'$ such that if $\langle r_i, c_j\rangle \in \mathcal{M}'$, then $\langle r_j, c_i\rangle \in \mathcal{M}'$. As noted above, the even length permutation cycles are split between two ordinary cycles in the bipartite graph when the symmetric edges are considered. By alternating the status of the edges according to the matching in one of the cycles, we can obtain length 2 permutation cycles. The decomposition of an even permutation cycle into length 2 permutation cycles is best seen in matrix terms. Consider the matching shown on the left below

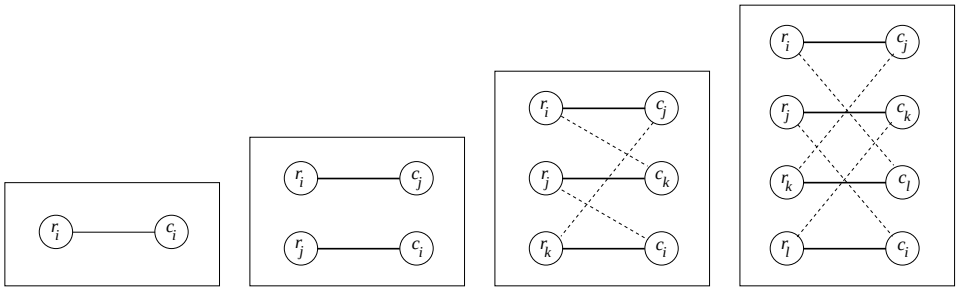

FIG. 2.3. *Permutation cycles of an automorphic matching. The matching edges are shown with bold solid lines. The other edges, shown with dashed lines, are present because of the symmetry of the matrix. Matching edges of the form* $\langle r_i, c_i\rangle$ *give a permutation cycle of length 1 (first subfigure); two matching edges of the form* $\big(\langle r_i, c_j\rangle, \langle r_j, c_i\rangle\big)$ *give a permutation cycle of length 2 (second subfigure); length 3 (third subfigure) and length 4 (fourth subfigure) permutation cycles are also shown.*

$$\begin{array}{c} \\ 1\\2\\3\\4\\5\\6 \end{array}\!\!\begin{array}{c} \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array}\\ \begin{pmatrix} & \times & & & & * \\ * & & \times & & & \\ & * & & \times & & \\ & & * & & \times & \\ & & & * & & \times \\ \times & & & & * & \end{pmatrix} \end{array} \qquad \begin{array}{c} \\ 1\\2\\3\\4\\5\\6 \end{array}\!\!\begin{array}{c} \begin{array}{cccccc} 1 & 2 & 3 & 4 & 5 & 6 \end{array}\\ \begin{pmatrix} & \times & & & & * \\ \times & & * & & & \\ & * & & \times & & \\ & & \times & & * & \\ & & & * & & \times \\ * & & & & \times & \end{pmatrix} \end{array}. \tag{2.1}$$

In the matrix on the left of (2.1), the original matching whose entries are marked by $\times$ is automorphic and corresponds to a length 6 permutation cycle. Using the entries symmetric to the matching entries, shown with $*$, we obtain two ordinary cycles: first $\big(\langle r_1, c_2\rangle$,$(c_2, r_3)$,$\langle r_3, c_4\rangle$,$(c_4, r_5)$,$\langle r_5, c_6\rangle$, $(c_6, r_1)\big)$ and second $\big(\langle r_2, c_3\rangle$, (c_3, r_4), $\langle r_4, c_5\rangle$, (c_5, r_6), $\langle r_6, c_1\rangle$, $(c_1, r_2)\big)$ which share the original matching edges evenly. Now by taking the second cycle and alternating the status of the edges according to the matching, we obtain $\big((r_2, c_3)$, $\langle c_3, r_4\rangle$, (r_4, c_5), $\langle c_5, r_6\rangle$, (r_6, c_1), $\langle c_1, r_2\rangle\big)$. As seen, the new set of matching edges form three length 2 permutation cycles: $\big(\langle r_1, c_2\rangle, \langle r_2, c_1\rangle\big)$, $\big(\langle r_3, c_4\rangle, \langle r_4, c_3\rangle\big)$, and $\big(\langle r_5, c_6\rangle, \langle r_6, c_5\rangle\big)$. The resulting matching is shown on the right of (2.1). As is clear, the two matchings $\mathcal{M}$ and $\mathcal{M}'$ have the same cardinality; if $\mathcal{M}$ is of the maximum cardinality, so is $\mathcal{M}'$.

Recall from Section 2.1 that after the maximum matching on the sample matrix was modified to be automorphic, three more closed m-sequences are created: $CS_6 = \big(\langle r_{17}, c_{10}\rangle, \langle r_{10}, c_{17}\rangle\big)$, $CS_7 = \big(\langle r_{15}, c_{14}\rangle, \langle r_{14}, c_{15}\rangle\big)$, and $CS_8 = \big(\langle r_1, c_{11}\rangle, \langle r_{11}, c_1\rangle\big)$. Among all closed m-sequences only the fifth one $CS_5 = \big(\langle r_{13}, c_5\rangle, \langle r_5, c_{18}\rangle, \langle r_{18}, c_{13}\rangle\big)$ is an odd-length permutation cycle, hence does not undergo any transformation. The closed m-sequences CS_1, CS_2, CS_4, and CS_6 to CS_8 are length 2 permutation cycles, and they do not undergo any transformation. We are left with $CS_3 = \big(\langle r_{19}, c_4\rangle, \langle r_4, c_{12}\rangle, \langle r_{12}, c_{21}\rangle, \langle r_{21}, c_{19}\rangle\big)$. As CS_3 is a permutation cycle of length 4, it can be decomposed into two length 2 permutation cycles: $\big(\langle r_{19}, c_4\rangle, \langle r_4, c_{19}\rangle\big)$, and $\big(\langle r_{12}, c_{21}\rangle, \langle r_{21}, c_{12}\rangle\big)$. The resulting matching is shown in Fig. 2.4.

We note that the transformation described above is performed to facilitate the proofs that will follow in the next section. They modify the matching, but do not change the matched rows or columns contrary to the transformation discussed in the previous subsection.

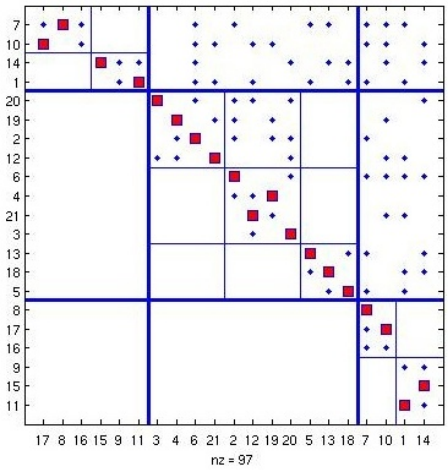

FIG. 2.4. *The automorphic maximum matching shown contains only length-2 or odd permutation cycles.*

3. The block triangular form of symmetric matrices. As shown in the previous section, any maximum matching in the bipartite graph of a structurally singular, symmetric matrix A can be transformed into an automorphic one consisting of odd permutation cycles and length 2 permutation cycles. Therefore, we assume that we have a matching $\mathcal{M}$ with those properties. We recall the uniqueness of the sets H_R, S_R, V_R, H_C, S_C, and V_C of the block triangular form of a matrix (Fact 1.2).

Before we prove our main theorem, we start with a series of lemmata.

LEMMA 3.1. *The odd permutation cycles are confined to only one block on the diagonal of the block triangular form.*

Proof. Permutation cycles of length 1 are trivially confined to only one block. Recall from Section 2.2 that an odd permutation cycle is a part of an ordinary cycle. Since the matching edges within the ordinary cycle are in the blocks on the diagonal of the btf, having those nonzeros straddle more than one block on the diagonal can only be possible if the submatrix confined to the rows and columns of the cycle is partly decomposable. However, as noted towards the end of Section 1.1, matrices whose sparsity structure corresponds to a superset of odd length cycles are fully indecomposable. □

We have a result which is stronger than the previous lemma.

LEMMA 3.2. *The odd permutation cycles are confined to the square block.*

Proof. From Lemma 3.1, an odd length permutation cycle is confined to a single block. Take an odd-length permutation cycle $\mathcal{C} = \big(\langle r_i, c_j\rangle, \langle r_j, \cdot\rangle, \ldots, \langle \cdot, c_i\rangle\big)$ and suppose, for the sake of contradiction, that $\mathcal{C}$ is in the horizontal block. Since all row vertices in $\mathcal{C}$ are in H_R, each one of these vertices reaches an unmatched column with an alternating path. Suppose row r_i reaches, with an alternating path, an unmatched column c_u, i.e., $r_i \xrightarrow{\mathcal{M}} c_u$ without going through other vertices in the permutation cycle $\mathcal{C}$ (this assumption is not weaker but gives a cleaner argument). Let $P = \big((r_i, c_l), \langle c_l, r_{i+1}\rangle, \ldots, \langle \cdot, r_{i+t}\rangle, (r_{i+t}, c_u)\big)$ be that alternating path. Due to the symmetry of the matrix, the same path exist in the reverse direction from row r_u to column c_i. That is we have the path $P^T = \big((r_u, c_{i+t}), (c_{i+t}, \cdot), \ldots, (c_{i+1}, r_l), (r_l, c_i)\big)$ in the bipartite graph of the matrix.

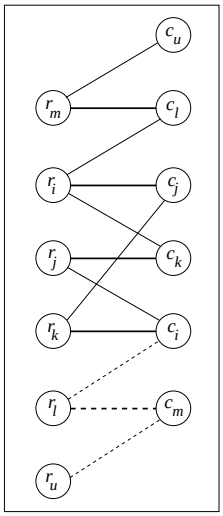

FIG. 3.1. *An example for the proof of Lemma 3.2. The solid bold lines correspond to edges of a maximum matching. A length 3 permutation cycle* $\mathcal{C} = \big(\langle r_i, c_j\rangle, \langle r_j, c_k\rangle, \langle r_k, c_i\rangle\big)$ *is shown; the solid lines represent the edges symmetric to that of the permutation cycle. Column* c_u *is not matched and reachable from row* r_i *with the alternating path* $P = \big((r_i, c_l), \langle c_l, r_m\rangle, (r_m, c_u)\big)$. *The symmetric path* P^T *is shown with dashed lines. Row* r_u *is not matched, as the matching is automorphic. Column* c_m *should have a mate, and since it reaches an unmatched row, it should be in the vertical block. It is shown in the proof that the column* c_m *should have been matched to* r_l*; this matching edge is shown with a dashed bold line. Therefore, the path* $\big((r_u, c_m), \langle c_m, r_l\rangle, (r_l, c_i), \langle c_i, r_k\rangle, (r_k, c_j), \langle c_j, r_i\rangle, (r_i, c_l), \langle c_l, r_m\rangle, (r_m, c_u)\big)$ *is an augmenting path, contradicting the fact that* $\mathcal{M}$ *is a maximum matching.*

We now investigate the status of the edges in P^T with respect to the given maximum matching $\mathcal{M}$. We show that P^T must be an alternating path. In other words, it must have the form $P^T = \big((r_u, c_{i+t}), \langle c_{i+t}, \cdot\rangle, \ldots, \langle c_{i+1}, r_l\rangle, (r_l, c_i)\big)$. Note that since r_u and c_u are both unmatched, P^T being an alternating path implies that the path $\big((r_u, c_{i+t}), c_{i+t} \xrightarrow{\mathcal{M}} c_i \xrightarrow{\mathcal{M}} r_i \xrightarrow{\mathcal{M}} r_{i+t}, (r_{i+t}, c_u)\big)$ is an augmenting path, contradicting the assumption that $\mathcal{M}$ is a maximum matching. We first note that $c_i \xrightarrow{\mathcal{M}} r_i$ as c_i and r_i are in an odd permutation cycle. Note that since $\mathcal{M}$ is automorphic, row r_u is not matched. Consider the last row vertex r_{i+t} in P. Column c_{i+t} is the first column vertex in P^T. Since r_u is not matched, c_{i+t} should have a mate (otherwise $\mathcal{M}$ would not be a maximum matching). Therefore, c_{i+t} is in the vertical block (being a matched vertex reaching an unmatched row). Now, since r_{i+t} is in the horizontal block and c_{i+t} is in the vertical one, due to Lemma 3.1 the vertices r_{i+t} and c_{i+t} cannot be in an odd permutation cycle. Therefore, they are in a length 2 permutation cycle. That is, if $\text{mate}(r_{i+t}) = c_x$, then $\text{mate}(c_{i+t}) = r_x$. Consider the next column vertex c_{i+t-1} in P^T. It should have a mate, otherwise an augmenting path $r_u \xrightarrow{\mathcal{M}} c_{i+t-1}$ exists, and should be in the vertical block. With the same reasoning as above, it is matched to the row that corresponds to the mate of row r_{i+t-1} in P. Therefore, the path P^T is an alternating path symmetric to P. Figure 3.1 displays the arguments for a length 3 permutation cycle $\mathcal{C} = \big(\langle r_i, c_j\rangle, \langle r_j, c_k\rangle, \langle r_k, c_i\rangle\big)$ and an alternating path $P = \big((r_i, c_l), \langle c_l, r_m\rangle, (r_m, c_u)\big)$.

With similar arguments, it can be shown that odd permutation cycles cannot be confined to the vertical block. Therefore, the odd permutation cycles are confined to

the square block. ∎

COROLLARY 3.3. *For each $\langle r_i, c_j \rangle \in \mathcal{M}$ in the horizontal block, we have $\langle r_j, c_i \rangle \in \mathcal{M}$. Similarly, for each $\langle r_k, c_l \rangle \in \mathcal{M}$ in the vertical block, we have $\langle r_l, c_k \rangle \in \mathcal{M}$.*

We have a refinement of the previous corollary.

LEMMA 3.4. *The length 2 permutation cycles are not contained entirely in the horizontal or vertical blocks.*

Proof. We prove the lemma for the horizontal block, that is we show that length 2 permutation cycles are not contained in the horizontal block; the vertical block case is similar. Suppose, for the sake of contradiction, $\langle r_i, c_j \rangle \in \mathcal{M}$ and its symmetric counterpart $\langle r_j, c_i \rangle \in \mathcal{M}$ are in the horizontal block. As in the proof of Lemma 3.2, we take an unmatched column c_u that is reachable from row r_i with an alternating path. Again, due to $\mathcal{M}$ being automorphic, row r_u is not matched. However, as in the proof of Lemma 3.2, we have an alternating path from column c_i to unmatched row r_u, contradicting the fact that c_i is in the horizontal block. ∎

We are now ready to state and prove the following theorem regarding the block triangular form of a structurally rank deficient symmetric matrix.

THEOREM 3.5. *Given a structurally rank deficient symmetric matrix A, let H_R, S_R, V_R, H_C, S_C, V_C be the sets in the block triangular form of A. Then, the set of indices of rows in H_R is equal to the set of indices of columns in V_C; the set of indices of the rows in S_R is equal to the set of indices of the columns in S_C; the set of indices of the rows in V_R is equal to the set of indices of the columns in H_C.*

Proof. Since A is square and structurally rank deficient, we know that both horizontal and vertical blocks are present in the block triangular form. The square block may be missing.

Consider a matching edge $\langle r_i, c_j \rangle$ in the horizontal block. As shown in Lemma 3.2, it is not in an odd permutation cycle and hence, as noted in Corollary 3.3, $\langle r_j, c_i \rangle \in \mathcal{M}$. We know from Lemma 3.4 that $\langle r_j, c_i \rangle \in \mathcal{M}$ is not in the horizontal block. Two cases remain to be investigated: $\langle r_j, c_i \rangle$ is either in the square block or in the vertical block. For the sake of contradiction, suppose $\langle r_j, c_i \rangle$ is in the square block. As in the proof of Lemma 3.2, we take an unmatched column c_u that is reachable from row r_i with an alternating path. Again, due to $\mathcal{M}$ being automorphic, row r_u is not matched. However, as in the proof of Lemma 3.2, we have an alternating path from column c_i to the unmatched row r_u, contradicting the fact that c_i is in the square block. Similar arguments can be used to show that for a $\langle r_k, c_l \rangle$ in the vertical block, $\langle r_l, c_k \rangle$ is in the horizontal block. Therefore, a matching edge $\langle r_i, c_j \rangle$ is in the horizontal block if and only if the matching edge $\langle r_j, c_i \rangle$ is in the vertical block.

We have established two results. First, the set of indices of rows in H_R is equal to the set of indices of columns in V_C. Second, the set of column indices that are matched to the rows in H_R is equal to the set of row indices that are matched to the columns in V_C. These equivalence relations are shown in Fig. 3.2. As the matching is automorphic, the set of indices of the unmatched columns (K in the figure) is equal to the set of indices of the unmatched rows (L in the figure); we have thus established the equivalence between the sets of indices of the columns in H_C and the rows in V_R. Since the matrix is square, the set of indices of the of remaining rows S_R is equal to the set of indices of the remaining columns S_C. ∎

Once the equivalence relations stated in Theorem 3.5 are established, it is easy to recover a structural symmetry in the block triangular form.

COROLLARY 3.6. *The block triangular form of a structurally singular, symmetric matrix can be permuted to be symmetric around the anti-diagonal.*

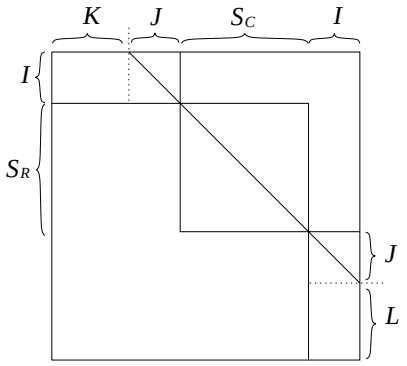

FIG. 3.2. *A matching is shown by the slanted line; the rows in I are matched to columns in J, and the columns in I are matched to rows in J. The borders of the horizontal, square, and vertical blocks are shown with solid lines. The dashed lines divide the set of columns in the horizontal block and the set of rows in the vertical block into two sets. The index sets K and L are equal as the set of indices of unmatched columns is equal to the set of indices of unmatched rows.*

This can be achieved by fixing a permutation of the rows in H_R, S_R, and V_R, and then by reorganizing V_C, S_C and H_C such that the reverse order within these later blocks match those of H_R, S_R, and V_R, respectively. It is possible to refine this form by looking at the fine structure of the square block A_S. We first need the following theorem.

THEOREM 3.7. *Let R_i and C_i be the set of row and column indices of the ith irreducible block in the fine decomposition of the square block A_S. Then, either $R_i = C_i$; or $R_i \cap C_i = \emptyset$ and there exists an irreducible block j in the fine decomposition of A_S with $R_j = C_i$ and $C_j = R_i$.*

Proof. Take the ith irreducible block and suppose for the sake of contradiction $R_i \neq C_i$ and $R_i \cap C_i \neq \emptyset$. Define three sets: first $I = R_i \cap C_i$, second $I_1 = R_i \setminus I$, and third $J_2 = C_i \setminus I$. With this partitioning of the rows and columns, we can permute the ith irreducible block into the following form

$$A_i = \begin{array}{c} \\ I \\ I_1 \end{array} \begin{array}{c} \begin{array}{cc} I & J_2 \end{array} \\ \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix} \end{array}. \tag{3.1}$$

First note that $A_{12} \neq O$ and $A_{21} \neq O$, otherwise the block will be reducible. Now consider the larger square submatrix consisting of the set of row indices $I \cup I_1 \cup J_2$ and the same set of column indices

$$A_L = \begin{array}{c} \\ I \\ I_1 \\ J_2 \end{array} \begin{array}{c} \begin{array}{ccc} I & J_2 & I_1 \end{array} \\ \begin{pmatrix} A_{11} & A_{12} & A_{21}^T \\ A_{21} & A_{22} & * \\ A_{12}^T & * & A_{22}^T \end{pmatrix} \end{array}. \tag{3.2}$$

For the columns whose indices are in the set I_1 and the rows whose indices are in the set J_2 to be in a different block from i, the submatrix of A_L defined by the row indices $I \cup J_2$ and the column indices $I \cup I_1$ should be reducible. But the matrix

$$\begin{array}{c} \\ I \\ J_2 \end{array} \begin{array}{c} \begin{array}{cc} I & I_1 \end{array} \\ \begin{pmatrix} A_{11} & A_{21}^T \\ A_{12}^T & A_{22}^T \end{pmatrix} \end{array} \tag{3.3}$$

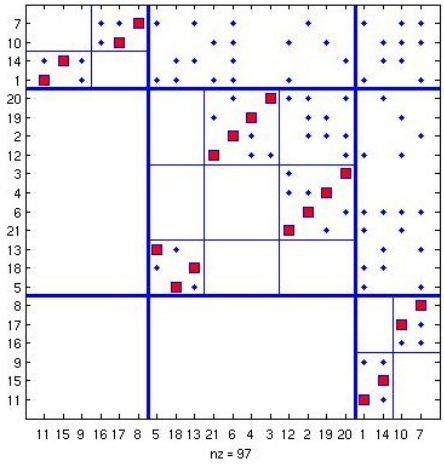

FIG. 3.3. *The matrix permuted blockwise so that the resulting matrix is symmetric along the anti-diagonal.*

is the transpose of A_i (see (3.1)), as $A_{11} = A_{11}^T$. Since A_i is irreducible, so is its transpose shown in (3.3). Therefore, the rows whose indices are in the set J_2 and the columns whose indices are in the set I_1 cannot be in a different block from the one that contains I.

We have established that either $R_i = C_i$ or $R_i \cap C_i = \emptyset$. If $R_i \cap C_i = \emptyset$, all matching nonzeros in this block should be in length 2 permutation cycles, as the odd permutation cycles are irreducible. Therefore, we have another irreducible block j with $R_j = C_i$ and $C_j = R_i$ and the proof is completed. □

Having defined the fine structure of the square block A_S, we refine Corollary 3.6 by using that structure. As before, let R_i and C_i denote the set of row and column indices of the ith irreducible block of A_S. We will order the rows of A and then apply that order in the reverse direction to the columns. We first order the rows in the horizontal block. Then, we order the rows in square block A_S using the fine structure as follows. Let i be an irreducible block whose rows are yet to be ordered. If the set of column indices C_i is equal to the set of row indices R_i, order R_i. If $C_i \cap R_i = \emptyset$, order R_i and then the rows that are symmetric counterparts of the columns whose indices are in the set C_i. After this blockwise ordering of all the rows in the square block A_S, we order the rows in the vertical block. We do not specify the order of the rows in a subblock—it can be arbitrary. Now applying the order obtained for the rows to the columns in the reverse direction results in a matrix that is symmetric along the anti-diagonal. Figure 3.3 shows the matrix of the previous figures permuted to be symmetric around the anti-diagonal. We emphasize that the matching itself is not symmetric around the anti-diagonal because of the odd permutation cycle on indices 5, 13, and 18. The matching entries can be placed to be symmetric around the anti-diagonal, but then the symmetry in other parts of the matrix can be lost.

Notice that one can also permute the matrix so that the resulting matrix is symmetric along the main diagonal; this time by ordering the rows as described above and then by applying the same permutation to the columns. The resulting permuted matrix is shown in left side of Fig. 3.4. One can further permute the matrix such that the square block of the Dulmage-Mendelsohn decomposition comes

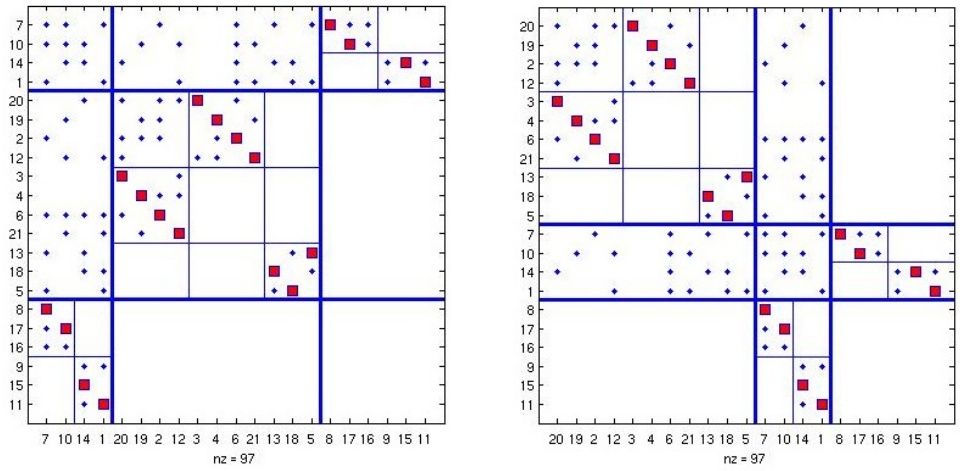

FIG. 3.4. *The matrix permuted blockwise so that the resulting matrix is symmetric along the main diagonal. When the square block of the coarse decomposition is permuted before the other blocks, a block tridiagonal form is obtain using the three blocks of the coarse Dulmage-Mendelsohn decomposition, with the last block on the diagonal being void of nonzeros*

before the rectangular ones as shown in the right side of the same figure. With this ordering the coarse Dulmage-Mendelsohn decomposition can be seen as inducing a block tridiagonal form with the last block on the diagonal being empty.

We note that Theorem 3.5, Corollary 3.6, and Theorem 3.7 hold for structurally full rank, symmetric matrices. In particular, the fine decomposition of such a matrix has square irreducible blocks with a row index set R_i and a column index set C_i where either $R_i = C_i$ or there exists another irreducible block j (with row index set R_j and column index set C_j) with $R_j = C_i$ and $C_j = R_i$.

4. Concluding remarks. We have presented observations on the block triangular form of structurally symmetric matrices. The two main results are that (i) if the matrix is structurally rank deficient, the horizontal and vertical blocks are transposes of each other, (ii) each block in the square submatrix has the same set of row and column indices or those two sets are totally disjoint and there is another square block being the transpose of the first one. These symmetrical structures help us recover the symmetry of the original matrix in the Dulmage-Mendelsohn decomposition (or in the block triangular form).

We note, for historical reasons, that Dulmage and Mendelsohn note (at the last paragraph in their seminal paper [7]) that they think it may be worthwhile to study the decomposition for symmetric bipartite graphs. We think that the observations in this paper could address some of their concerns.

Acknowledgements. We thank two anonymous referees and Cleve Ashcraft, also a referee, for their constructive comments and suggestions. Cleve also provided us with some Matlab code for the figures. We also thank Alex Pothen for his remarks on an earlier version of the paper and bringing the Mendelsohn-Dulmage theorem, restated in Section 2.1, to our attention.

REFERENCES

[1] C. Berge, *Two theorems in graph theory*, Proceedings of the National Academy of Sciences of the USA, 43 (1957), pp. 842–844.

[2] P. J. Cameron, *Notes on algebraic structures*. Available at http://www.intute.ac.uk/, last accessed 12 Feb, 2008.

[3] I. S. Duff, *On permutations to block triangular form*, Journal of the Institute of Mathematics and its Applications, 19 (1977), pp. 339–342.

[4] I. S. Duff, A. M. Erisman, and J. K. Reid, *Direct Methods for Sparse Matrices*, Oxford University Press, London, 1986. In preprint of second edition. To appear.

[5] I. S. Duff and S. Pralet, *Strategies for scaling and pivoting for sparse symmetric indefinite problems*, SIAM Journal on Matrix Analysis and Applications, 27 (2005), pp. 313–340.

[6] I. S. Duff and J. K. Reid, *An implementation of Tarjan's algorithm for the block triangularization of a matrix*, ACM Transactions on Mathematical Software, 4 (1978), pp. 137–147.

[7] A. L. Dulmage and N. S. Mendelsohn, *Coverings of bipartite graphs*, Canadian Journal of Mathematics, 10 (1958), pp. 517–534.

[8] ———, *Two algorithms for bipartite graphs*, SIAM Journal on Applied Mathematics, 11 (1963), pp. 183–194.

[9] E. Lawler, *Combinatorial Optimization: Networks and Matroids*, Dover, (Unabridged reprint of Combinatorial Optimization: Networks and Matroids, published by Holt, Rinehart, and Wilson, New York 1976), New York, 2001.

[10] L. Lovasz and M. D. Plummer, *Matching Theory*, North-Holland Mathematics Studies, Elsevier Science Publishers, Amsterdam, Netherlands, 1986.

[11] M. Marcus and H. Minc, *A Survey of Matrix Theory and Matrix Inequalities*, Dover, (Unabridged, unaltered republication of the corrected (1969) printing of the work published by Prindle, Weber, & Schmidt, Boston, 1964), New York, 1992.

[12] N. S. Mendelsohn and A. L. Dulmage, *Some generalizations of the problem of distinct representatives*, Canadian Journal of Mathematics, 10 (1958), pp. 230–241.

[13] A. Pothen, *Sparse null bases and marriage theorems*, PhD thesis, Department of Computer Science, Cornell University, Ithaca, New York, 1984.

[14] 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.