

Prime Factorization Theory of Networks¹

Shuo-Yen Robert Li, The Chinese University of Hong Kong

Guangyue Han, The University of Hong Kong

Yixian Yang, Beijing University of Posts and Telecommunications

¹ The work of the first two authors is partially supported by a grant from the University Grants Committee of the Hong Kong Special Administrative Region, China (Project No. AoE/E-02/08). The second author also gratefully acknowledges the support of Research Grants Council of the Hong Kong Special Administrative Region, China under grant No HKU 701708P.

Contents

Preface.....	3
Chapter 1. Matching Theory.....	5
Section 1.1. Basic terminology and notation.....	5
Section 1.2. The Edmonds matching algorithm	7
Section 1.3. Prime factorization of networks with respect to matching.....	17
Chapter 2. Prime Factorization of Networks	32
Section 2.1. Definitions and notation	34
Section 2.2. Equivalence between templates.....	39
Section 2.3. Related work.....	42
Chapter 3. Prime Factorization with Respect to the Template Δ_2	44
Section 3.1. An Edmonds-type algorithm	45
Section 3.2. Prime factorization of networks with respect to Δ_2	55
Chapter 4. Prime Factorization with Respect to the Template Δ_n	69
Section 4.1. An Edmonds-type algorithm	70
Section 4.2. Prime factorization of networks with respect to Δ_n	76
Chapter 5. Prime Factorization with Respect to the Template X_n	89
Section 5.1. An Edmonds-type algorithm	90
Section 5.2. Prime factorization of networks with respect to X_n	109

Preface

One way to factorize an integer is via the recursive process of identifying a factor at a time. A polynomial, or more generally, an element in a unique factorization domain can also be factorized into primes, which cannot be further factorized. Bearing the same spirit, many other types of mathematical objects can also be recursively factorized in similar fashions. Examples include Abelian groups, stationary Markov chains and invariant measures. The characteristic of the *prime factorization* for each type of objects depends on their algebraic structure. For instance, it yields the product of *prime* factors as well as a *unit* factor in the case of a unique factorization domain.

A “network” means a set of nodes interconnected by links. An exemplifying network consists of service centers interconnected by channels for the *interflow* traffic of service requests. In the abstract form, a network is called a *graph* and the links interconnecting the nodes are called *edges*. A *network partition* is to classify the vertices into classes according to a given *template*, an algorithmic approach to find a *maximum network partition* naturally leads to a *network factorization*, where a graph is decomposed into *prime* pieces through removing a *factorizer*. Since a graph by itself lacks the necessary algebraic structure to create such a sense of prime factorization, the notion of a “template” is coined so that network factorization is always with respect to a predetermined template. Different templates can lead to drastically different ways of factorization.

The classical *matching* is a special case of *network partition* and *network factorization*, although there is a fundamental difference between the viewpoints. A graph that does not possess a perfect matching is regarded as being “deficient” in the matching theory. Network partition and network factorization theory, on the other hand, treats such deficiency as “complexity.” The more

deficient a graph is, the higher the complexity. Thus network factorization decomposes a graph into subgraphs of minimal complexities. To motivate this new concept, Chapter 1 reviews the basic matching theory in the language of a special form of network factorization.

Then Chapter 2 sets up the notion of a template and the general concept of network factorization with respect to a template. At the same time, the variety of templates of interest is reduced through *equivalence* to just those templates X_n and Δ_n , where $n \geq 2$. Matching coincides with network factorization with respect to X_2 . The remaining chapters then deal with network factorization with respect to Δ_2 , X_n and Δ_n , where $n \geq 3$.

The prime factorization theory of networks traces back to an unpublished manuscript [37] of S.-Y. R. Li in 1978, which was intended for a paper. The theory grew in length over time, and a summary [38] was tentatively published in a conference in 1993. Y. X. Yang joined the effort in scrutinizing the technical detail during the early 1990's. The writing and publishing process however still lagged behind. A recent collaborative work with G. Han at the Institute of Network Coding of The Chinese University of Hong Kong finally brought the lengthy process to a closure.

Chapter 1. Matching Theory

Matching means making pairs among a group of objects. Every object in the set is matchable to some, but not necessarily all, other objects. The typical matching problem is to match as many pairs as possible. Matching problems arise in a wide variety of contexts, in both daily life and mathematical study. For instance, in the classical Marriage Problem, a girl is matchable to a boy in the same community if she knows the boy; and the problem asks whether all girls can be matched to different boys. If such a matching is not possible, then what is the maximum number of matched pairs and how to form such pairs algorithmically? These problems can always be cast in terms of graph theory, where an object is represented as a *vertex* in a *graph* and a matchable pair by an *edge*.

This chapter reviews the classical matching theory, which will be generalized to network factorization theory in subsequent chapters. In order to set up the terminology and background knowledge for the general network factorization theory, the presentation of the classical matching theory in this chapter deviates somewhat from the conventional approach in the literature.

Section 1.1. Basic terminology and notation

This section sets up some basic terminology and notation.

A *graph* G is a pair (V, E) , where V is a finite nonempty set and E is a family of two-element subsets of V . An element of V is called a *vertex* of the graph and hence V itself the *vertex set*. An element of E is called an *edge* of the graph and E itself the *edge set*. The *order* of a graph $G = (V, E)$ is defined as $|V|$, the cardinality of V .

An edge $e = (u, v)$ is said to *join* the two vertices u and v , and the two vertices, which are often referred to as the *endpoints* of e , are said to be *adjacent* to each other and are *incident to* this edge. Furthermore, when two edges are incident to a common vertex, they are said to be *adjacent edges*.

A graph $G_1 = (V_1, E_1)$ is *isomorphic* to a graph $G_2 = (V_2, E_2)$ if between them there exists an *isomorphism*, which means a one-to-one mapping from V_1 onto V_2 that preserves adjacency among vertices. It is easy to see that isomorphism is an equivalence relation.

A graph $G_2 = (V_2, E_2)$ is called a *subgraph* of a graph $G_1 = (V_1, E_1)$ when $V_2 \subset V_1$ and $E_2 \subset E_1$; alternatively, we say that G_1 is a *supergraph* of G_2 . The deletion of an edge subset E' from a graph $G = (V, E)$ yields the subgraph $G - E' = (V, E - E')$; in particular, the deletion of an edge e yields the subgraph $G - e = (V, E - \{e\})$. The deletion of a vertex subset V' from a graph $G = (V, E)$ yields the subgraph $G - V' = (V - V', E')$, where E' denotes the set E minus those edges incident to some vertex in V' ; in particular, the deletion of a vertex v yields the subgraph $G - v = (V - \{v\}, E')$, where E' means the set E minus those edges incident to v . The *induced* subgraph of $G = (V, E)$ on $V' \subset V$ means the graph (V', E') , where E' consists of those edges that are joining two vertices in V' . Similarly, the *induced* subgraph on $E' \subset E$ means the graph (V', E') , where V' consists of those vertices that are incident to at least one edge in E' .

The *degree* of a vertex v in a graph G is the number of edges that are incident to it and is denoted by $\deg_G(v)$ or simply $\deg(v)$ when G is clear from the context.

A *path* in a graph G is a sequence of edges $(u_0, v_0), (u_1, v_1), \dots, (u_n, v_n)$ such that $v_i = u_{i+1}$ for $i = 0, 1, \dots, n-1$, and all the vertices $u_0, u_1, \dots, u_n, v_n$ are distinct. We often denote such a path by $(u_0, u_1, \dots, u_n, v_n)$, the sequence of distinct vertices on the path; and we refer to such a path as an u_0 - v_0 path, and u_0, v_0 as the *terminal vertices* of this path. A *cycle* in a graph G is a sequence of edges $(u_0, v_0), (u_1, v_1), \dots, (u_n, v_n)$ such that $v_i = u_{i+1}$ for $i = 0, 1, \dots, n-1$ and $v_n = u_0$, and all the vertices u_0, u_1, \dots, u_n are distinct. We often denote such a cycle by (u_0, u_1, \dots, u_n) , the sequence of distinct vertices on the cycle. The *length* of a path (or cycle) is defined to be the number of edges on the path (or cycle).

Two distinct vertices u and v in a graph G are said to be *connected* to each other if there is a u - v path in G . The connectedness among vertices is an equivalence relation. It partitions the vertex set into equivalence classes. The subgraph induced on each equivalent class is called a *connected component* (or simply *component*) of the graph. A graph is said to be a *connected* or *disconnected* depending whether there is only one component or not.

An edge subset M is said to be a *matching* in G , if no two edges in M are incident to the same vertex. With respect to a given matching M , a vertex u is said to be *covered* if there is an edge in M incident to u ; otherwise, the vertex u is said to be *exposed*. A matching containing the maximum number of edges is called a *maximum matching*; the cardinality of a maximum matching is denoted by $\nu(G)$.

Two special types of graphs are of particular interest: A graph is said to be *bipartite* if its vertex set can be partitioned into two subsets S and T such that every edge of the graph has one endpoint in S and the other in T . A graph is said to be *complete* if all vertices are adjacent. A complete graph of order n is conventionally denoted by K_n .

Section 1.2. The Edmonds matching algorithm

Given a matching M in a graph $G = (V, E)$, an *alternating path (cycle)* is a path (cycle) whose edges are alternately in M and not in M . An *augmenting path* with respect to M is an alternating path between two exposed vertices. For two matchings M, N in G , let $M \oplus N$ denote the symmetric difference between M and N , that is, $M \oplus N = (M - N) \cup (N - M)$. The following lemma is straightforward.

Lemma 1.2.1. *Let M and N be two matchings of G . Then, every connected component of the subgraph of G induced on $M \oplus N$ takes one of the following forms (see Figure 1-1):*

- (a) *A cycle of even length whose edges are alternately in M and N .*

- (b) A path of even length whose edges are alternately in M and N .
- (c) A path of odd length whose edges are alternately in M and N and whose terminal vertices are both exposed by M .
- (d) A path of odd length whose edges are alternately in M and N and whose terminal vertices are both exposed by N .




Figure 1-1: Components of the symmetric difference $M \oplus N$.

The following theorem has been proven in [43].

Theorem 1.2.2. *M is a maximum matching in G if and only if G admits no augmenting path with respect to M.*

Proof. The “only if” part. Suppose that there is an augmenting path P with respect to M in G . Treat P as a set of edges. Then, $M \oplus P$ is a matching whose cardinality exceeds that of M . Thus M is not a

maximum matching.

The “if” part. Suppose that M is not a maximum matching. Let N be a matching with $|N| > |M|$. Consider the induced subgraph of G on $M \oplus N = (M - N) \cup (N - M)$. Since $|N| > |M|$, at least one connected component of this subgraph contains more edges from N than from M . From Lemma 1.2.1, every component of this subgraph is either a cycle or a path, whose edges are alternately in M and N . Moreover, a component with more edges from N than from M can only be a path of odd length whose terminal vertices are both exposed by N . This is an augmenting path with respect to M . ■




Figure 1-2: A bipartite graph G with a non-maximum matching M

Figure 1-2 depicts a bipartite graph G , of which the highlighted edges constitute a matching. This is not a maximum matching because of the presence of the augmenting path $(5, 8, 3, 9)$.

Recursive application of Theorem 1.2.2 yields the following well-known fact.

Theorem 1.2.3 (Mendelsohn-Dulmage Theorem). *All vertices covered by an arbitrary matching of a graph are also covered by some maximum matching.*

We shall describe an algorithm that determines whether a matching is maximum. For this purpose, we need the notion of graph *contraction*:

Definition 1.2.4. Given a vertex subset W of a graph $G = (V, E)$, the *contraction* of W into a new vertex w means a mapping from V to $(V \setminus W) \cup \{w\}$ that preserves $V \setminus W$ colluding W into w . The contraction naturally induces a *contracted graph* with the vertex set $(V \setminus W) \cup \{w\}$ so that the contraction preserves vertex adjacency.

Definition 1.2.5. Let M be a matching on a graph G and $(x_0, x_1, \dots, x_{2n})$ an odd-length cycle such that (x_{2k-1}, x_{2k}) is in M for $1 \leq k \leq n$. Let G' denote the graph obtained from contracting this cycle. The image M' of M under the contraction clearly forms a matching on G' , which is called the *induced matching* by M on G' .

Algorithm 1.2.6 (The Edmonds matching algorithm [12]). Given a matching M on G , this algorithm determines whether M is a maximum matching and, when M is not, finds an augmenting path with respect to M . Write $G_0 = G$ and $M_0 = M$. The algorithm will construct a sequence of graphs G_t , $0 \leq t \leq \tau$, and a matching M_t on each G_t . In the end, whether there is an augmenting path with respect to M_τ in G_τ will be apparent. If there is not, then M is a maximum matching on G . If there is, then, for every t , an augmenting path with respect to M_{t+1} in G_{t+1} induces an augmenting path with respect to M_t in G_t . The graph G_t will be associated with, besides the matching M_t , an acyclic subgraph T_t in which every vertex is labeled either *even* or *odd* so that T_t is a bipartite graph between even and odd vertices. Figure 1-3 illustrates G_t , M_t and T_t for a generic t .

Initially, let T_0 consist of z_1, z_2, \dots, z_d , all the vertices exposed by M . Label all these d vertices

even. Following the construction of G_t , M_t and T_t , the next iterative step in the algorithm, to be described shortly, shall achieve exactly one of the following:

- (a) Keep both G_t and M_t the same, whereas grow T_t by adding an odd vertex, an even vertex, and two edges. The first edge is between an existing even vertex and the new odd vertex; the second is between the new vertices and belongs to M_t . At the end of this step, increase the index t by 1.
- (b) Contract an odd cycle in T_t (and G_t) to obtain T_{t+1} (and G_{t+1}), and let M_t induce a matching M_t on G_{t+1} . At the end of this step, increase the index t by 1.
- (c) Identify an augmenting path of M_t , and recursively find an augmenting path with respect to M . The algorithm terminates, that is, t is the final index τ .
- (d) The algorithm terminates with the assertion of M being a maximum matching on G .

Given G_t , M_t and T_t , the next iterative step starts by looking for an edge of G_t that is

- not an edge of T_t ,
- incident to at least one even vertex of T_t , and
- not incident to any odd vertex of T_t .

The iterative step incurs the following separate cases:

Case 1. Such an edge does not exist. Then G does not admit an augmenting path with respect to M , so M is a maximum matching. The algorithm terminates. ((d) is achieved.)

Case 2. Such an edge exists. Let (e, f) be such an edge of G_t , where e is an even vertex of G_t .

Case 2.1. f is not a vertex of T_t . Find the unique $g \in V(G_t)$ such that (f, g) is in M_t (such g necessarily exists since M_t covers all vertices in $G_t - V(T_t)$). Then add the two vertices f, g and the two edges (e, f) , (f, g) into the graph T_t to obtain T_{t+1} . The vertex f is labeled odd and g even. Set $G_{t+1} = G_t$, $M_{t+1} = M_t$. Increase t by 1 and (a) is achieved (See the illustration of Figure 1-4.)

Case 2.2. f is an even vertex of T_t . Let $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n} = e)$ be the unique alternating path in T_t with respect to M_t with x_0 exposed, let $(y_0, y_1, y_2, \dots, y_{2m-1}, y_{2m} = f)$ be the unique alternating path in T_t with respect to M_t with y_0 exposed (necessarily, all $(x_{2i-1}, x_{2i}), (y_{2j-1}, y_{2j})$ are necessarily in M_t).

Case 2.2.1. $x_0 = y_0$. Then let $k \geq 0$ be the largest index with $x_k = y_k$ (necessarily, k must be an even integer). Thus $\{x_k, x_{k+1}, \dots, x_{2n} = e, y_{2m} = f, \dots, y_{k+1}, y_k = x_k\}$ is an odd cycle in the subgraph induced on $V(T'_t)$ with respect to M_t . Contract this cycle into a single vertex to obtain G_{t+1} , and set M_{t+1} to be the induced matching by M_t on G_{t+1} . Increase t by 1 and (b) is achieved (See the illustration of Figure 1-5.)

Case 2.2.2. $x_0 \neq y_0$. Then $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n} = e, y_{2m} = f, y_{2m-1}, \dots, y_1, y_0)$ is an augmenting path in G_t with respect to M_t . This constructs an augmenting path with respect to M in G by recursive invocation of Lemma 1.2.7 below. The algorithm terminates and (c) is achieved (See the illustration of Figure 1-6.)

In conclusion, M is a maximum matching if and only if Case 2.2.2 never occurs throughout the execution of the algorithm. ■




Figure 1-3: G_t , T_t , M_t are constructed in Algorithm 1.2.6 by time t . An even vertex of T_t is represented by a rectangle, an odd vertex of T_t by a hollow circle, and a vertex in $G_t - V(T_t)$ by a solid circle. An edge of G_t is regarded as outside T_t if it is incident to a vertex outside T_t . The matching M_t is indicated by highlighted edges. The figure also displays (inside rectangles) those groups of vertices in G that have been contracted into even vertices of T_t .




Figure 1-4: New vertices are added to T_t to obtain T_{t+1} (Case 2.1 in Algorithm 1.2.6).

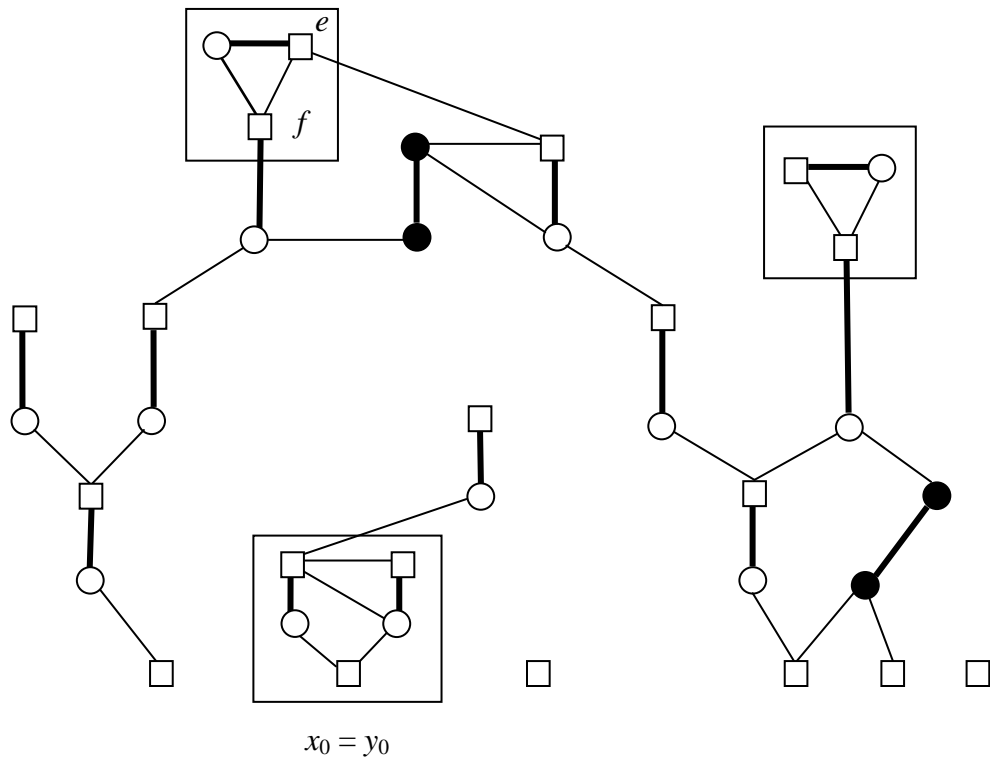


Figure 1-5: An odd cycle is contracted (Case 2.2.1 in Algorithm 1.2.6).

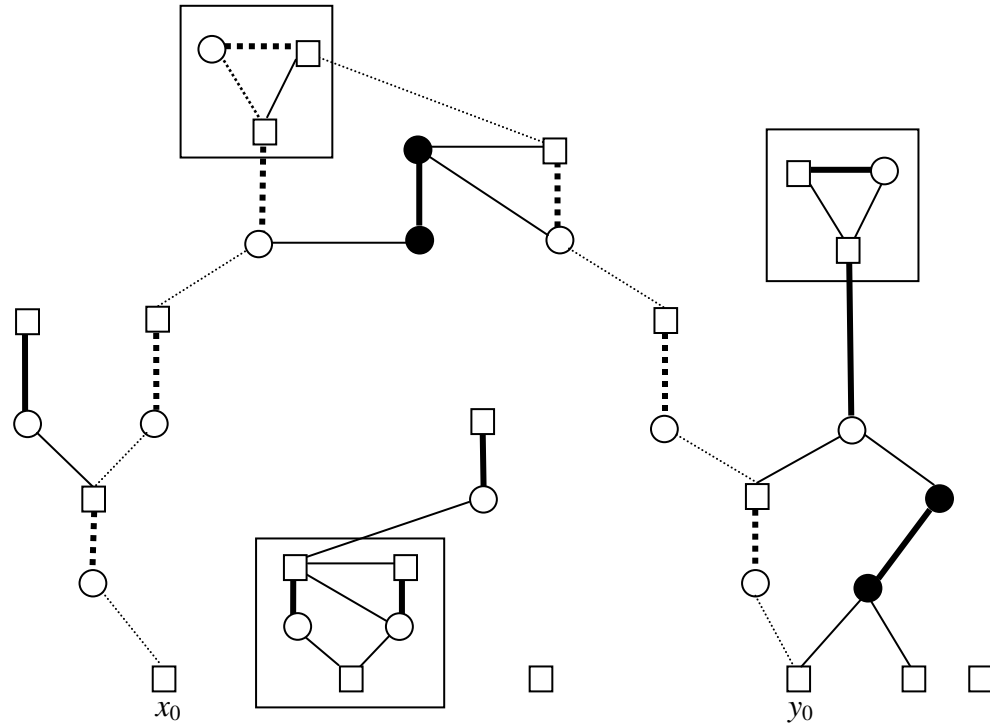


Figure 1-6: based on an augmenting path of G_t , an augmenting path (highlighted as dotted path) of G is found (Case 2.2.2 in Algorithm 1.2.6)

Lemma 1.2.7. *For any t , there is an augmenting path in G_t with respect to M_t if and only if there is an augmenting path in G_{t+1} with respect to M_{t+1} .*

Proof. In this proof, we use the notation adopted in the algorithm. Assume that G_{t+1} is obtained by contracting an odd cycle C in G_t into a new vertex w in G_{t+1} .

The “only if” part. Suppose that there is an augmenting path $P = (u_0, u_1, \dots, u_{2l+1})$ in G_t . We consider the case when P shares some common edges with C (otherwise, the proof is trivial). Now, traverse P from u_0 to u_{2l} . Let u_i (u_j) be the first (last) vertex on P that is also on C . Then either (u_{i-1}, u_i) or (u_j, u_{j+1}) is not matched. Without loss of generality, assume that (u_{i-1}, u_i) is not matched.

Then $(u_0, u_1, \dots, u_i, w, x_{k-1}, x_{k-2}, \dots, x_0)$ is an augmenting path in G_{t+1} .

The “if” part. Suppose that there is an augmenting path $P = (u_0, u_1, \dots, u_{2l+1})$ in G_{t+1} . We shall assume that P passes through w because the opposite case is trivial. Let $u_i = w$. Without loss of generality, assume that (u_{i-1}, u_i) is matched, and thus (u_i, u_{i+1}) is not. Let (v, u_{i+2}) be a pre-image of (u_{i+1}, u_{i+2}) under the contraction mapping. One then checks that from u_k to v , there is always an alternating path P_1 of even length consisting of only edges in C . Concatenating (u_0, u_1, \dots, u_i) , P_1 and $(v, u_{i+2}, \dots, u_{2l+1})$ gives us an augmenting path in G_t . ■

We are now ready for justification of Algorithm 1.2.6. First, suppose that Case 2.2.2 does occur, that is, an augmenting path is found in T_τ . Then, by Lemma 1.2.7, there is an augmenting path in G and hence M is not a maximum matching. Now, suppose that Case 2.2.2 never occurs throughout the execution of Algorithm 1.2.6. One then checks that each T_i consists of d connected components, each of which contains exactly one exposed vertex by M . This further implies that between any two of the exposed vertices, there is no augmenting path in T_τ , and thus there is no augmenting path in G_τ . Repeatedly applying Lemma 1.2.7 for all t , we then conclude that there is no augmenting path in $G_0 = G$, so M is a maximum matching.

Remark 1.2.8. Algorithm 1.2.6 identifies an augmenting path with respect to any non-maximum matching. Often there are multiple choices for the augmenting path in each step. By selecting the augmenting path in a strategic way, the computational complexity in finding a maximum matching can be contained to $O(N^{2.5})$, where N is the number of vertices (See [15], [42]). ■

Section 1.3. Prime factorization of networks with respect to matching

This section recasts the classical matching theory using the language of prime factorization theory of networks.

Definition 1.3.1. A graph is said to be *regular* when it allows a perfect matching and otherwise

singular. The number of vertices exposed by a maximum matching of a graph G is called the *dimension* of G and denoted by $\dim(G)$.

A graph with a positive dimension is one without a perfect matching. The dimension of a graph is also referred to as the *deficiency* by some authors (See [31], for example.) In the present theory of network factorization, we shall treat this notion as a measure of how “elementary” the graph is. The theory, in fact, will “partition” a graph of a large dimension into subgraphs of smaller dimensions.

Apparently $\dim(G)$ shares the same parity as $|G|$. Thus, for every vertex x in a graph,

$$(1.3-1) \quad \dim(G-x) = \dim(G) \pm 1$$

A vertex x is called a *pole* when $\dim(G-x) = \dim(G) - 1$ or, equivalently, when x is exposed by a maximum matching. Otherwise, x is called a *zero*. A zero that is adjacent to at least one pole is called a *root*. It then follows from (1.3-1) that, for every vertex subset S ,

$$(1.3-2) \quad -|S| \leq \dim(G-S) - \dim(G) \leq |S|$$

Lemma 1.3.2. *If S is a vertex subset of a graph G such that*

$$(1.3-3) \quad \dim(G-S) = \dim(G) + |S|$$

then S consists of only zeroes.

Proof. For any vertex x in S ,

$$\begin{aligned} \dim(G-x) + |S \setminus x| &\geq \dim((G-x) - (S \setminus x)), \text{ by (1.3-2)} \\ &= \dim(G-S) \\ &= \dim(G) + |S|, \text{ by (1.3-3)} \\ &= \dim(G) + 1 + |S \setminus x|. \end{aligned}$$

Thus $\dim(G-x) \geq \dim(G) + 1$ and hence x is a zero. ■

Apparently, the dimension of a disconnected graph is equal to the sum of the dimensions of its components. In view of (1.3-2), the number of singular components in the graph $G-S$ is at most

$\dim(G)+|S|$.

Definition 1.3.3. A vertex subset S of G is called a *factorizer* if the number of singular components in the graph $G-S$ is exactly $\dim(G)+|S|$.

Remark 1.3.4. Let S be a factorizer of G . Then, by (1.3-2), all singular components of $G-S$ are by themselves graphs with dimension 1 and hence $\dim(G-S) = \dim(G)+|S|$. Therefore:

- Every vertex in S is a zero of G by Lemma 1.3.2.
- Given a matching, a vertex in S can be matched to at most one exposed vertex from $G-S$. Thus, a maximum matching must match every vertex in S to some vertex from a distinct singular component of $G-S$. This is illustrated by highlighted edges of a graph G with $\dim(G) = 1$ in Figure 1-7.

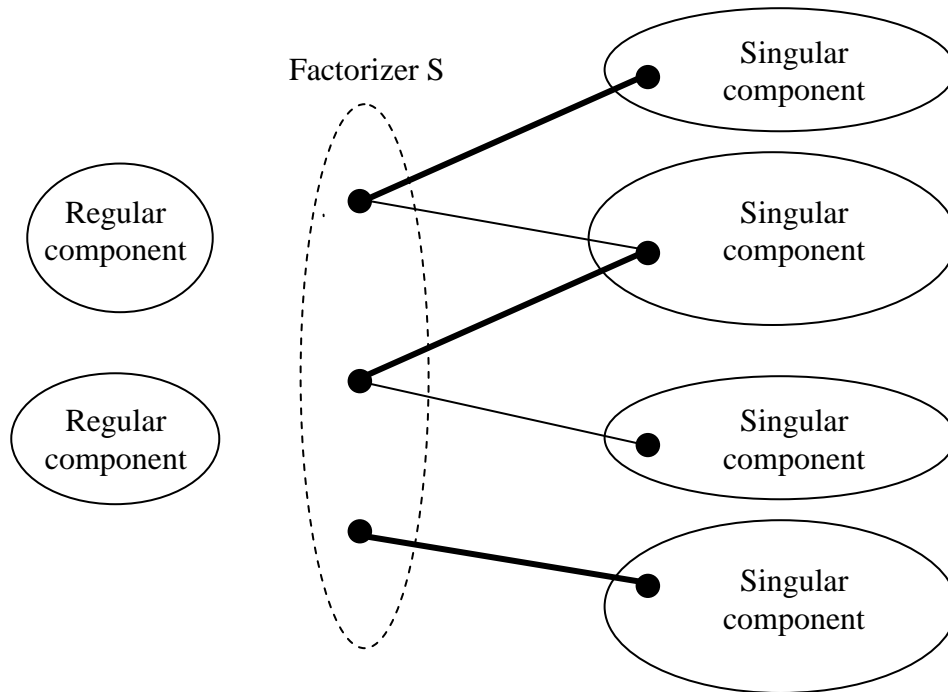


Figure 1-7: Removal of a factorizer S from G with a maximum matching.

The following lemma follows from the definition of factorizer.

Lemma 1.3.5. *If S is a factorizer of G and S' is a factorizer of $G-S$, then $S \cup S'$ is a factorizer of G .*

Definition 1.3.6. A connected graph with no non-empty factorizer is called a *prime* graph. A factorizer S of a graph G is said to be *primary* when all singular components of $G-S$ are prime (a regular graph is not prime, since every single vertex in any regular graph is a factorizer); furthermore S is said to be *prime* if all components of $G-S$ are prime.

If the removal of a factorizer does not make all components prime, then the non-prime components can be further factorized. Repeatedly applying this factorization process, we would eventually reach a stage where all the remaining components are prime graphs. Then, by Lemma 1.3.5, all vertices that have been removed during the process constitute a prime factorizer. Thus every graph possesses at least one prime factorizer, which may possibly be the empty set.

We next present the Gallai-Edmonds structure theorem, the fundamental theorem in the classical matching theory. We prove this theorem by telescoping the recursive invocations of Algorithm 1.2.6. First, we need the definition of blossom and related lemmas.

Definition 1.3.7. A *blossom* is a special graph defined recursively as follows. A single-vertex graph is a blossom. If a graph G contains a cycle of an odd length such that the contraction of this cycle yields a blossom, then G is also a blossom.

Lemma 1.3.8. *A blossom is a prime graph with dimension 1. Moreover, every vertex in a blossom is a pole.*

Proof. To prove an arbitrary blossom B is a prime graph, by Lemma 1.3.2, it suffices to prove the non-existence of zeroes in any blossom B . The proof is by induction on $|B|$. Let C be a cycle of odd length in B whose contraction into x transforms B into a blossom B' of a smaller order. By induction, the graph B' contains no zeroes and $\dim(B') = 1$. Thus every vertex of B' is a pole of B' (remember that any vertex is either a pole or a zero).

For any vertex y in C , let M_y be the maximum matching of C isolating only y . Now consider a maximum matching M_x of B' isolating only x . Then one readily checks that $M_x \cup M_y$ is a maximum matching of B isolating only y , which implies y is a pole of B .

For any vertex z in B but not in C , let M_z be a maximum matching of B' isolating only z . Let w be the vertex in C that is matched by M_z , and let M_w be the maximum matching of C isolating only w . Then one readily checks that $M_z \cup M_w$ is a maximum matching on B isolating only z , which implies that z is a pole. Thus every vertex in B is a pole. ■

Lemma 1.3.9. *Let M be a maximum matching on a blossom B with exposed vertex x . Then, from x to any other vertex in B , there always exists an alternating path of even length with respect to M .*

Proof. Note that for any vertex y in B , there is a unique maximum matching isolating only y . Let N be the maximum matching exposing y . Then x, y must be in the same component of $M \oplus N$, taking the form of $(x_0 = x, x_1, \dots, x_{2n} = y)$ such that the edges (x_k, x_{k+1}) , $0 \leq k \leq 2n-1$, are alternately in N and M (see Lemma 1.2.1). This is an alternating path of an even length from x to y . ■

We also need the following lemma.

Lemma 1.3.10. *Let M be a maximum matching on a graph G . If an alternating path with respect to M starts at an exposed vertex of M , then the vertices on the path are alternately poles and roots.*

Proof. Let $(x_0, x_1, \dots, x_{2n})$ be such an alternating path with respect to M , where x_0 is a vertex exposed by M . Then, all (x_{2k+1}, x_{2k+2}) , $0 \leq k \leq n-1$, must be in M . For any x_{2k} , one checks that $M_k = M \oplus (x_0, x_1, \dots, x_{2k})$ is a maximum matching of G , isolating x_{2k} , which implies x_{2k} is a pole of G .

We next prove that each x_{2k+1} is a root. It suffices to prove that x_{2k+1} is not a pole. Assuming that x_{2k+1} is a pole, we shall derive a contradiction. Let N be a maximum matching that exposes x_{2k+1} . From Lemma 1.3.2, every component of $M_k \oplus N$ is either a path or a cycle whose edges are alternately in M_k and N . Since M_k and N are both maximum, there is no augmenting path with

respect to either of them according to Theorem 1.2.2. Hence every component of $M_k \oplus N$ is of an even length. Since x_{2k} is exposed by M_k , the component containing x_{2k} can only be an even length path. Let this component take the form of $P = (y_0 = x_{2k}, y_1, y_2, \dots, y_{2l})$, where all (y_{2j}, y_{2j+1}) are in N and (y_{2j+1}, y_{2j+2}) are in M_k . One then checks that $(N \oplus P) \cup (x_{2k}, x_{2k+1})$ is a matching with larger cardinality than N , which is a contradiction. ■

Following [38], we now state and prove the Gallai-Edmonds structure theorem using the language of prime factorization theory of networks.

Theorem 1.3.11 (The Gallai-Edmonds structure theorem). *Let P denote the set of poles in a graph G and R the set of roots. Then,*

- (a) $G - (P \cup R)$ is a regular graph, on which every maximum matching of G induces a perfect matching.
- (b) Every connected component of the induced subgraph on P is a blossom. Moreover, every vertex in R is adjacent to vertices in at least two such blossoms.
- (c) R is a primary factorizer of G .
- (d) Let F be the induced subgraph of G on $P \cup R$. Then every vertex in P (resp. R) is a pole (resp. root) of the graph F . Moreover, $\dim(G) = \dim(F)$.
- (e) $\nu(G) = (|V| - c(P) + |R|)/2$, where $c(P)$ denotes the number of odd components of the induced subgraph on P .

Proof. If G is a regular graph, then any vertex in G is a zero, thus the theorem trivially follows. In the remaining proof, we only consider the case when G is singular, namely, there exists at least one pole in G .

Let M be a maximum matching on G , and let z_1, z_2, \dots, z_d denote the vertices exposed by M ,

and apply Algorithm 1.2.6 on G with respect to M . It can be easily checked that, throughout the iterative process, the following 5 basic properties are satisfied:

- (1) Every odd vertex in T_t is a single vertex of the original graph G , so is every vertex in $G_t - V(T_t)$. Every even vertex in T_t is a contracted blossom of G .
- (2) If (f, g) is in M_t , then either both f and g or neither of them are vertices in $G_t - V(T_t)$.
- (3) Every odd vertex in T_t is adjacent to exactly two even vertices and is covered by M_t .
- (4) Every connected component of T_t contains exactly one exposed vertex by M .
- (5) The number of even vertices in T_t exceeds the number of odd vertices by exactly d .

We deduce from the above five properties the sixth property:

- (6) In the original graph G , for any exposed (by M) vertex z_j , any odd vertex v in T_t , any vertex w labeled even in T_t , there exist an alternating path of odd length with respect to M from z_j to v , and an alternating path of even-length with respect to M from z_j to any vertex in the blossom corresponding to w (by Property (1), w is a contracted blossom of G).

We first prove the “odd vertex” part of (6). By property (4), there exists a unique alternating path of odd length in T_t with respect to M_t that connects z_j to v . Let this path be $(x_0 = z_j, x_1, x_2, \dots, x_{2n-1}, x_{2n}, x_{2n+1} = v)$, where each (x_{2i-1}, x_{2i}) is in M_t for $1 \leq i \leq n$. For $0 \leq i \leq n$, let B_{2i} be the blossom in G corresponding to the even vertex x_{2i} , and p_{2i} be the vertex in B_{2i} that is adjacent to x_{2i-1} (thus p_{2i} is exposed by M on B_{2i}), and q_{2i} be a vertex in B_{2i} that is adjacent to the odd vertex x_{2i+1} . By Lemma 1.3.9, there exists an alternating path of even length in G with respect to M from p_{2i} and q_{2i} , $0 \leq i \leq n$. These alternating paths and the paths $(q_{2i}, x_{2i+1}, p_{2i+2})$, $0 \leq i < n$, and (q_{2n}, x_{2n+1}) can be concatenated to form an alternating path in G with respect to M of odd length from z_j to v . A similar argument can be applied to prove the “even vertex” part.

Since M is a maximum matching, Case 2.2.1 never occurs during the execution of Algorithm

1.2.6. It follows from Property (6) that M_t is a maximum matching on G_t for all t . Observe that M_t exposes exactly d vertices in T_t and none in $G_t - V(T_t)$. So, we have

$$(7) \dim(G_t) = d \text{ for all } t.$$

Assume that the algorithm terminates after the τ -th step and outputs G_τ, M_τ, T_τ . Let R'_τ denote the set of odd vertices in T_τ . Let e be any even vertex in T_τ . By Properties (4) and (3), there exists an alternating path of even length with respect to M_τ in T_τ that connects an exposed vertex z_j to e . By Lemma 1.3.10, we conclude that

$$(8) \text{ Every even vertex in } T_\tau \text{ is a pole of } G_\tau.$$

By Property (1), every vertex in R'_τ is also a vertex of the original graph G , and by Property (5), there are exactly $d + |R'_\tau|$ blossom components in $G - V(R'_\tau)$. Note that even vertices are only adjacent to odd vertices in G_τ , so in the original graph G , these blossom components are only adjacent to vertices in R'_τ . It then follows that $\dim(G - x) \geq d$ for any vertex x in $G_\tau - V(T_\tau)$. Since M induces a perfect matching on $G_\tau - V(T_\tau)$, we conclude that $\dim(G - x) > d$ for any vertex x in $G_\tau - V(T_\tau)$. We therefore reach the following:

$$(9) \text{ Any vertex in } G_\tau - V(T_\tau) \text{ is not a pole of } G.$$

$$(10) \quad R'_\tau \text{ is a factorizer of } G.$$

From property (10) and Lemma 1.3.2, we know every vertex of R'_τ is a zero of G . This, together with property (9), proves the “only if” part of the following property.

$$(11) \quad \text{A vertex of } G \text{ is a pole if and only if it belongs to a blossom in } G \text{ that is contracted into an even vertex of } T_t.$$

To prove the “if” part of property (11): Let B be a blossom in G that is contracted into an even vertex e' of T_τ . We shall prove that any vertex e in B is a pole of G . It follows from Properties (7)

and (8) that $\dim(G_n - e') = d - 1$. Since the dimension of a graph is unchanged when a blossom is contracted into a single vertex, we have $\dim(G - B) = \dim(G_n - e')$. Thus, $\dim(G - e) \leq \dim(G - B) + \dim(B - e) = d - 1 + 0 = \dim(G) - 1$. This completes the proof of property (11).

By Properties (10) and (1) as well as Lemma 1.3.2, every odd vertex is a zero of G . Thus, every vertex in R'_τ is a root of G , by Properties (3) and (11). On the other hand, by Property (11), there are no roots other than the odd vertices. Therefore, we have proved that

(12) A vertex of G is a root if and only if it is an odd vertex of T_τ , that is, $R'_\tau = R$.

Obviously M covers all the vertices in $G - V(T_\tau)$, and there are no edges in M joining a vertex in $V(T_\tau)$ and a vertex in $G - V(T_\tau)$ (otherwise, the edge would be added to T_τ). So we prove that

(13) M induces a perfect matching on $G - V(T_\tau)$, which implies $G - V(T_\tau)$ is a regular graph.

Statement (a) follows from Properties (11), (12) and (13), while (b) follows from Properties (11) and (13). From Properties (10) and (12), we know that R is a factorizer of G . Furthermore, by Properties (7), (1) and (13), every singular connected component of $G - V(R)$ is a blossom (which is a prime graph, by Lemma 1.3.8). Thus, (c) is also proved.

We next prove (d). From (a), the graph $G - V(T_\tau)$ is regular. Thus, $d = \dim(G) \leq \dim(G - V(T_\tau)) + \dim(T_\tau) = \dim(T_\tau)$. On the other hand, M induce a matching on T_τ which exposes d vertices, that is, $\dim(T_\tau) \leq d$. This establishes the following property:

(14) $\dim(G) = \dim(T_\tau) = d$.

Let x be a vertex in P , that is, $\dim(G - x) = d - 1$. Assume that M' is a maximum matching on G which exposes the vertex x . From Property (12), there exists no poles of G in $G - F$, thus M_n induces a matching on F which exposes x and some other $d - 1$ vertices. Therefore, $\dim(T_{\tau - x}) \leq \dim(G - x) = d - 1 = \dim(T_\tau) - 1$, which implies that x is also a pole of T_τ , namely,

(15) every vertex in P is a pole of T_τ .

From (a) and (b), R is also a factorizer of F . Thus every vertex in R is a zero of F , by Lemma 1.3.2.

By Property (3), we prove that:

(16) Every vertex in R is a root of F .

Properties (14), (15) and (16) together establish (d).

Finally, it follows from statement (d) and Properties (3), (4), (5) that $\dim(G) = \dim(F) = c(P) - |R|$. We then have statement (e):

$$v(G) = (|V(G)| - \dim(G))/2 = (|V(G)| - c(P) + |R|)/2. \blacksquare$$

Example 1.3.12. Let G be the graph in Figure 1-8. Then the following statements hold.

- (1) G is a singular graph with $\dim(G) = 1$.
- (2) The vertex set $P = \{1, 2, 3, 4, 5, 6\}$ is the set of poles in G .
- (3) The vertex set $\{a, b, c, d, e\}$ is the set of zeros in G .
- (4) The subgraphs induced on $\{2, 3, 6\}$ and $\{2, 3, 4, 5, 6\}$ are two blossoms.
- (5) The vertex set $R = \{a\}$ is the only root of G which forms a primary factorizer.
- (6) The vertex sets $\{a, b, d\}$ and $\{a, c, d\}$ are the only two prime factorizers of G .
- (7) The subgraph $G - (P \cup R)$ is shown in Figure 5.9 (a) which is a regular graph.
- (8) The induced subgraph on P is shown in Figure 1-9 (b). Both of its connected components are blossoms of orders 1 and 5, respectively. The induced subgraph F on $P \cup C$ is shown in Figure 1-9 (c). It is clear that $\dim(F) = \dim(G) = 1$ and the set $\{a\}$ is the only root of F and all the other vertices in F are poles. \blacksquare




Figure 1-8: An example

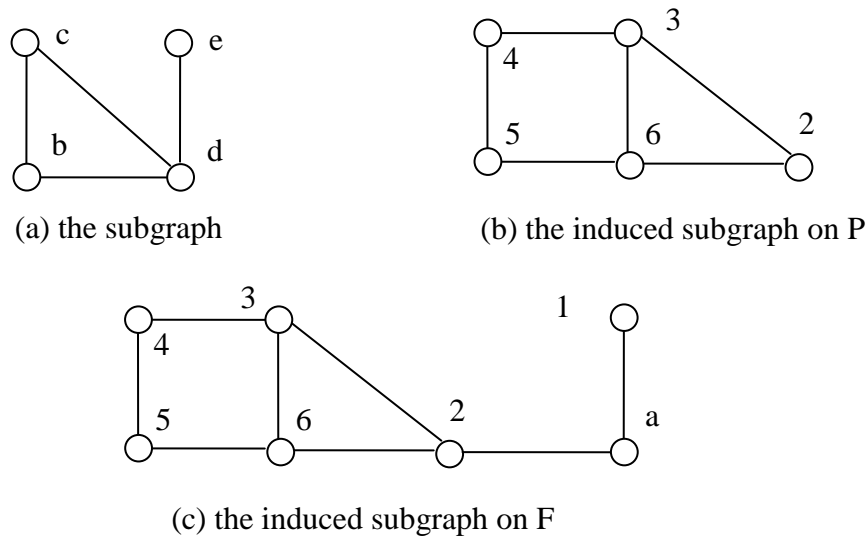


Figure 1-9: Subgraph and induced subgraph in the example path

The following theorem, which is due to Berge [6], naturally follows from Theorem 1.3.11.

Theorem 1.3.13. (Berge Formula). *For any graph G , $\dim(G) = \max\{c(G-X)-|X|: X \subset V\}$, where $c(G-X)$ denotes the number of odd components in the subgraph $G-X$.*

Proof. We first prove that $\dim(G) \geq \max\{c(G-X)-|X|: X \subset V\}$. Let M be a maximum matching of G . Let X be any vertex subset, and let G_1, \dots, G_k , $k := c(G-X)$, denote all the odd components of $G-X$. Among these components, renumbering if necessary, let G_1, \dots, G_j be those containing a

vertex exposed by M . Then for each $j+1 \leq i \leq k$, there is at least one edge in M from X to G_i , which implies $|X| \geq k-j$. On the other hand, $\dim(G) \geq j$ since each of G_1, \dots, G_j contains an exposed vertex. Hence $\dim(G) \geq j \geq k-|X| = c(G-X)-|X|$. We then conclude that $\dim(G) \geq \max\{c(G-X)-|X|: X \subset V\}$. On the other hand, if we choose X to be R , then by Theorem 1.3.11 (statement e)), we have $c(G-R)-|R| = |V|-2v(G) = \dim(G)$, which establishes $\dim(G) = \max\{c(G-X)-|X|: X \subset V\}$. ■

The following theorem is an immediate corollary of Theorem 1.3.13.

Theorem 1.3.14. (Tutte's Theorem). *A graph $G = (V, E)$ has a perfect matching if and only if for every vertex subset $S \subset V$, $c(G-S) \leq |S|$, where $c(G-S)$ denotes the number of odd components in $G-S$.*

The following theorem characterizes all the prime graphs.

Theorem 1.3.15. *The following statements are equivalent for a graph G .*

- (a) G is a prime graph.
- (b) G is a prime graph with dimension 1.
- (c) G is connected, and all the vertices of G are poles.
- (d) G is a blossom.

Proof. (a) \Rightarrow (b). First note that G must be singular, since otherwise any vertex in G constitutes a factorizer, which contradicts that assumption that G is prime. So, we only need to prove that the dimension of G is not strictly greater than 1. Suppose that, by contradictions, G has dimension strictly greater than 1. Let M be a maximum matching of G with exposed vertices z_1, z_2, \dots, z_d , where $d > 1$. Apply Algorithm 1.2.6 to G with respect to M . Then, there are odd vertices in T_π , and all the odd vertices constitute a factorizer by Statement (c) of Theorem 1.3.11, a contradiction.

(b) \Rightarrow (c). Let M be a maximum matching of G with exposed vertices z_1 . Apply Algorithm 1.2.6 to

G with respect to M . Then there are no odd vertices in T_τ , thus T_τ consists of only one even vertex, which implies that all vertices are poles.

(c) \Rightarrow (d). By Theorem 1.3.11, all poles, thus all vertices, in G are contained in blossoms, which will be contracted into even vertices when Algorithm 1.2.6 terminates. Since there are no roots in G , all the vertices constitute exactly one blossom.

(d) \Rightarrow (a). This follows directly from Lemma 1.3.8. ■

Theorem 1.3.16. *Every primary factorizer of a graph contains all roots. Thus, the set of roots is the unique minimal primary factorizer.*

Proof. Let a maximum matching M of a graph G expose the vertices z_1, z_2, \dots, z_d . Apply Algorithm 1.2.6 to G with respect to M . Then, the roots of G are simply odd vertices in T_n . Let R denote the set of roots. Assume for contradiction that there is a primary factorizer S that does not contain R as a subset. Pick an odd vertex $x \in R \setminus S$. One then checks that there are more even vertices than odd vertices in the component of $G_n - S$ containing x . Since even vertices in T_τ correspond to blossoms in the original graph G , the component C of $G - S$ containing x is not a blossom. So, by Theorem 1.3.15, C is not prime. Apparently, C is a singular component. This contradicts the assumption that S is a primary factorizer. ■

Let the set of these vertices that are neither poles nor roots be called the *neutral factor* of the graph G . Then, by Theorem 1.3.11, the induced subgraph on the neutral factor is a regular graph; and if the neutral factor is deleted from the graph, the poles would remain poles, the roots would remain roots, and the dimension of the graph is unchanged.

Any maximum matching M can be decomposed into a perfect matching on the subgraph induced on the neutral factor and a maximum matching on the induced subgraph on poles and roots. For each of those blossoms that are connected components of the induced subgraph on poles, M

must match all but one of its vertices into pairs. One checks that M induces a maximum matching on the bipartite graph between (those poles not matched by M with other poles) and roots (see Figure 1-10 for an example, where matching is indicated by highlighted edges).

Thus every maximum matching on the original graph can be constructed by the following steps: First, construct a perfect matching on the subgraph induced on the neutral factor. Then, construct a maximum matching on the induced subgraph on poles. This will isolate exactly one vertex on each connected component. Finally, construct a maximum matching on the bipartite graph between those not yet matched poles and roots.

Example 1.3.17. For illustrative purposes, we shall characterize all maximum matchings of an example graph $G = (E, V)$ in Figure 1-10. Identification of poles, zeroes, roots, infinities, and blossoms can be done following Algorithm 1.2.6. We label the poles, zeroes, roots, and infinities in this graph using “ p ”, “ 0 ”, “ c ”, and “ i ”, respectively, and we further indicate roots by squares. Vertices labeled with “ $0i$ ” constitute the neutral factor, since they are neither poles nor roots. Encircled connected components of the induced subgraph on poles are all blossoms. Any maximum matching M can be decomposed into a perfect matching on the subgraph induced on the vertices labeled as “ $0i$ ” and a maximum matching on the induced subgraph on vertices labeled as “ p ”, “ pi ” or “ $0r$ ” (see Figure 1-10 for an example).




Figure 1-10: An example with a maximum matching

By Theorem 1.3.11, the set of all roots is a primary factorizer. Every prime factorizer consists of all roots plus some “0i” vertices whose removal would leave the remaining graph a disjoint union of blossoms. For the particular graph in Figure 1-10, there are six prime factorizers:

$\{c, d, e, f, v, y\}$, $\{c, d, e, f, v, z\}$, $\{c, d, e, f, u, x, y\}$, $\{c, d, e, f, u, x, z\}$, $\{c, d, e, f, w, x, y\}$ and $\{c, d, e, f, w, x, z\}$.

Chapter 2. Network Partition and Network Factorization

Many types of services are rendered by multiple service centers at different locations. Examples include airline reservation, ticketing, distributed computing, telephone operator service and switching, the Internet connection, etc. Service requests at one service center can be redirected to another provided that there is an *interflow* connection between the two centers. There are several possible reasons for interflow. One purpose of interflow is to alleviate traffic congestion at a single service center. This, at the same time, serves as a contingency measure against facility breakdown at any location. Also, some special service demands can be redirected to the designated service centre, e.g., a gateway for long-distance connection.

Between any two service centers, the feasibility of installing an interflow connection between them is often dictated by geographic, economic and political factors. The configuration of feasible connections thus defines a *network*, that is, a graph, wherein every *vertex* represents a service center and every *edge* a feasible interflow connection. The service network should be partitioned into interflow regions so that, when one service center in the region is in operation, all others in the same region can redirect traffic to it (this way, only one active center is needed in each geographic region during the light traffic time). Meanwhile, connectivity requirements impose restrictions on the topology of an interflow region. For one thing, the region must be a connected one in some sense. Also, there may be a limit on the size of a region or on the number of links to a vertex. The restrictions on the topology define a family of allowable shapes of an interflow region. For instance, if the topology of a region can only be either a single vertex or two adjacent vertices, then a partition of the graph into regions simply means a matching.

In partitioning a network into interflow regions subject to these restrictions, we try to avoid single-vertex regions, which represent service centers without any interflow connection. The

network partition problem is to partition the network into regions under the restricted topology and minimize the number of exposed vertices. Naturally, the optimal partition may not be unique.

One of the aims of *network factorization* theory is to give a simple characterization of all optimal partitions (as exemplified at the end of Chapter 1 when network partition is simply matching) by “labeling” all the vertices according to their “roles” in optimal partitions. With such a characterization, a network planner can then select among all optimal partitions to suit ad hoc considerations in individual applications.

Another aim of network factorization theory is to find ways to decompose a possibly complicated network into simpler “prime” subnetworks, which are typically easier to characterize. Such decomposition is in fact a conventional and powerful approach in many disciplines of mathematics: to analyze mathematical objects prohibitively complex, we often “decompose” them into smaller or simpler “pieces”, so that the study of the original objects can be reduced to that of smaller or simpler pieces. Prominent examples include:

- In number theory, the fundamental theorem of arithmetic states that any positive integer greater than 1 can be “uniquely” (up to some permutation) written as a product of prime numbers. A generalized version of the fundamental theorem of arithmetic in commutative ring theory states that every nonzero nonunit element in a unique factorization domain can be uniquely written as a product of prime elements.
- In algebra, the fundamental theorem of finite abelian groups states that every finite abelian group G can be expressed as a direct sum of cyclic subgroups of prime-power order. This fundamental theorem can be generalized to the case when the abelian group has zero rank and is finitely generated.
- In probability theory, all states of a stationary Markov chain can be classified into

disjoint classes, on each of which the original Markov chain induces an irreducible Markov chain. More generally, ergodic decomposition theorem, in loose terms, states that an invariant measure can be decomposed to a convex sum of ergodic measures.

Bearing the same spirit, for a given network partition, the proposed prime network factorization in this work factorizes a possibly complicated network into prime graphs, which, to some extent, can be characterized more explicitly.

Section 2.1. Definitions and notation

This section formulates network factorization theory in graph theoretic terms. Certain elementary results are derived for later use. We first provide definitions and notation for the factorization theory with respect to a *template*, which means a family of shapes under some simple restrictions. We then establish the equivalence relationship among different templates and singles out two sequences of templates that play the center roles of the theory.




Figure 2-1: (a) A template to be named X_4 in the sequel. (b) An X_4 -partition of a graph G , where intra-class edges under the partition are highlighted.

A *shape* means a graph up to isomorphism. A family of shapes is said to be *hereditary* if, for every member G of the family and every vertex x of G , every connected component of $G-x$ is isomorphic to a shape in the family (Recall from Section 1.1 that $G-x$ denotes the subgraph of G induced on all vertices other than x). A *template* means a family of connected shapes. Given a graph and a template Γ , a Γ -*partition* of the graph means a classification of its vertices into classes such that the induced subgraph on each class is isomorphic to a member of Γ .

Naturally every hereditary template Γ includes the shape of a single vertex. Under a Γ -partition of a graph, a vertex is said to be *exposed* if it forms a singleton class by itself. Meanwhile, a vertex that is not exposed is said to be *covered* by that partition. The Γ -*dimension* of a graph G means the

minimum number of vertices exposed by a Γ -partition on G and is denoted by $\dim(G, \Gamma)$. If a particular Γ -partition defines exactly $\dim(G, \Gamma)$ singletons, that partition is called a *maximum* Γ -partition on G . When $\dim(G, \Gamma) = 0$, the graph G is said to be Γ -regular and otherwise Γ -singular. For a Γ -regular graph, a maximum Γ -partition does not expose any vertex and is therefore called a *perfect* Γ -partition.

Example 2.1.1. Figure 2-1(a) shows a 4-member template Γ , and Figure 2-1(b) depicts a maximum Γ -partition of a graph G . Isolating the vertex t , this maximum Γ -partition in Figure 2-1(b) is not perfect. Thus the graph G is Γ -singular. In fact $\dim(G, \Gamma) = 1$. ■

Let Γ be a finite and hereditary template. The Γ -order of a vertex x in a graph G is defined as $\dim(G-x, \Gamma) - \dim(G, \Gamma)$. Obviously,

$$(2.1-1) \quad \dim(G-x, \Gamma) + 1 \geq \dim(G, \Gamma),$$

This implies that the Γ -order of a vertex is always greater than or equal to -1 . A vertex with the Γ -order equal to -1 is called a Γ -pole in G . On the other hand, define the *maximum* Γ -order of a graph G as

$$(2.1-2) \quad \Phi(G) = \max_{x \in V} \{ \dim(G-x, \Gamma) - \dim(G, \Gamma) \}$$

Let $\Phi(\Gamma)$ be the maximum Γ -order among for all shapes in Γ ($\Phi(\Gamma)$ is well-defined since Γ is finite). Since Γ is hereditary, $\Phi(\Gamma)$ is equal to the maximum number of neighboring vertices of a vertex in any shape in Γ that are pairwise non-adjacent. A vertex in G with the Γ -order equal to $\Phi(\Gamma)$ is called a Γ -zero. For later use, we also define a Γ -root in a graph as a vertex that is not a Γ -pole but is adjacent to at least one Γ -pole. As a counterpart to a Γ -root, a Γ -infinity is a vertex such that all adjacent vertices (if any) are Γ -zeroes.

A necessary and sufficient condition for a vertex to be a Γ -pole is it being exposed by a

maximum Γ -partition. Recall that for a vertex subset S in a graph G , the subgraph of G induced on all vertices outside S is denoted as $G-S$. By induction on $|S|$, inequalities (2.1-1) and (2.1-2) can be easily generalized to

$$(2.1-3) \quad \dim(G, \Gamma) - |S| \leq \dim(G-S, \Gamma) \leq \dim(G, \Gamma) + \Phi(\Gamma)|S|$$

Since all members of Γ are connected shapes, the Γ -dimension of a disconnected graph is equal to the sum of the Γ -dimension of its *connected components*. In view of the last inequality (2.1-3), there can be at most $\dim(G, \Gamma) + \Phi(\Gamma)|S|$ connected components of $G-S$ that are Γ -singular graphs.

A vertex subset S is called a Γ -factorizer of G if there are $\dim(G, \Gamma) + \Phi(\Gamma)|S|$ connected components of $G-S$ that are Γ -singular graphs (this implies that exactly $\dim(G, \Gamma) + \Phi(\Gamma)|S|$ connected components of $G-S$ have the Γ -dimension 1 and all others are Γ -regular graphs). A Γ -prime graph is defined as a connected graph that has no non-empty Γ -factorizer. A Γ -singular Γ -prime graph is called a Γ -blossom. A Γ -factorizer S is called a *primary* Γ -factorizer if all Γ -singular components of $G-S$ are Γ -blossoms (while the Γ -regular components may or may not be Γ -prime). A Γ -factorizer S of a graph is called a *prime* Γ -factorizer if all connected components of $G-S$ are Γ -prime graphs.

Example 2.1.2. Let Γ be the template in Figure 2-1(a) and G the graph in Figure 2-1(b). Clearly, vertices p, q, r , and t are Γ -poles, $\Phi(\Gamma) = 3$, and the vertex s is a Γ -zero. Thus vertices p, q, r , and t are all Γ -infinities of G , while vertex s is a Γ -root. There are four Γ -singular connected components of the subgraph $G-s$, namely the singletons $\{p\}$, $\{q\}$, $\{r\}$ and $\{t\}$, which naturally are Γ -blossoms. Since $\dim(G, \Gamma) + \Phi(\Gamma) = 4$, the set $\{s\}$ is a primary Γ -factorizer of G . Let $S = \{s, v\}$. There are exactly $\dim(G, \Gamma) + \Phi(\Gamma)|S| = 7$ connected components in $G-S$ and all of them are Γ -prime. Thus S is a prime Γ -factorizer of G . ■

Lemma 2.1.3. *Let S be a vertex subset of a graph G such that*

$$(2.1-4) \quad \dim(G-S, \Gamma) = \dim(G, \Gamma) + \Phi(\Gamma)/|S|,$$

then S only consists of Γ -zeroes. This is true in particular when S is a Γ -factorizer of G .

Proof. For any vertex x in S ,

$$\begin{aligned} \dim(G-x, \Gamma) + \Phi(\Gamma)/|S \setminus x| &\geq \dim((G-x)-(S \setminus x), \Gamma), && \text{by (2.1-3)} \\ &= \dim(G-S, \Gamma) \\ &= \dim(G, \Gamma) + \Phi(\Gamma)/|S|, && \text{by (2.1-4)} \\ &= \dim(G, \Gamma) + \Phi(\Gamma) + \Phi(\Gamma)/|S \setminus x| \end{aligned}$$

Thus $\dim(G-x, \Gamma) \geq \dim(G, \Gamma) + \Phi(\Gamma)$ and hence x is a Γ -zero by the definition of $\Phi(\Gamma)$. ■

Lemma 2.1.4. *Let S be a Γ -factorizer of a graph G . Then, every Γ -pole of G is a Γ -pole of $G-S$.*

Proof. Let x be a Γ -pole of G . By Lemma 2.1.3, x does not belong to S . By (2.1-3) we have

$$\begin{aligned} \dim(G-(S \setminus x), \Gamma) &\leq \dim(G-x, \Gamma) + \Phi(\Gamma)/|S| \\ &= \dim(G, \Gamma) - 1 + \Phi(\Gamma)/|S| \\ &= \dim(G-S, \Gamma) - 1 \end{aligned}$$

Thus x is a Γ -pole of $G-S$. ■

The following lemma follows directly from the definition of a Γ -factorizer.

Lemma 2.1.5. *If S is a Γ -factorizer of G and S' is a Γ -factorizer of $G-S$, then $S \cup S'$ is a Γ -factorizer of G .*

The implication of Lemma 2.1.5 is as follows: If a Γ -factorizer does not factor the graph into Γ -prime pieces, then it can be “enlarged”. Through iterations of such enlargement, we would eventually arrive at a prime Γ -factorizer, therefore we conclude that every graph possesses at least one prime Γ -factorizer, which may possibly be the empty set. Note that, in general, the prime Γ -factorizer may not be unique.

Section 2.2. Equivalence between templates

If two templates Γ_1 and Γ_2 are such that $\dim(G, \Gamma_1) = \dim(G, \Gamma_2)$ for all graphs G , we say that the two families Γ_1 and Γ_2 are *equivalent* to each other and write $\Gamma_1 \approx \Gamma_2$. A necessary and sufficient condition for the equivalence is that every shape in Γ_1 other than a single vertex is a Γ_2 -regular graph and vice versa. The equivalence implies $\Phi(\Gamma_1) = \Phi(\Gamma_2)$. It also implies that the notions of Γ_1 -pole, -zero, -root, -infinity, -factorizer, etc. are the same as their Γ_2 -counterparts. Thus, equivalent templates are interchangeable as far as network factorization theory is concerned.

Recall from the beginning of the chapter various networks of service centers via interflow connections. When Γ is one of the following special templates, Γ -partition of a network of service centers is of interest.

- (1) *Connected_n*: the family of connected graphs of order up to n .
- (2) *Central_n*: the family of *centrally connected* graphs, that is, a graph with one vertex adjacent to all other vertices, of order up to n .
- (3) *Complete_n*: the family of complete graphs of order up to n .
- (4) *Tree_n*: the family of *trees*, that is, connected graphs containing no cycles, of order up to n .

Under a *Connected_n*-partition of the network, customer traffic toward all vertices in a class can be served by any single vertex through interflow connections. Under a *Central_n*-partition, the central vertex in a class can receive traffic interflow *directly* from other vertices. With a *Complete_n*-partition, every vertex in the class can serve for this purpose. Meanwhile, a *Tree_n*-partition of a network requires connectedness among a class but disallows loops.

The next theorem asserts the equivalence of every hereditary template to one in the following two sequences:

- (5) X_n : the family of shapes $Star_k$, $1 \leq k \leq n$, where $Star_k$ is the centrally connected tree of order

k as illustrated in Figure 2-1(a).

(6) $\Delta_n : \Delta_n = X_n \cup \{K_3\}$, where K_3 represents the complete graph of order 3.

Note that $\text{Connected}_2 = \text{Central}_2 = \text{Complete}_2 = \text{Tree}_2 = X_2$, $\text{Connected}_3 = \text{Central}_3 = \Delta_3$, $\text{Complete}_3 = \Delta_2$, and $\text{Tree}_3 = X_3$.

Theorem 2.2.1. *Every finite and hereditary template Γ is equivalent to either $\Delta_{\Phi(\Gamma)+1}$ or $X_{\Phi(\Gamma)+1}$ depending upon whether the shape K_3 is a member.*

Proof. Since Γ is hereditary, the shape $\text{Star}_{\Phi(\Gamma)+1}$ is a member of Γ and hence so are all other members of $X_{\Phi(\Gamma)+1}$. Thus every multi-vertex shape in $X_{\Phi(\Gamma)+1}$ is a Γ -regular graph. Conversely, it suffices to prove that, for any multi-vertex shape G in Γ , G is $\Delta_{\Phi(\Gamma)+1}$ -regular.

We next show that there exists an induced subgraph H of G such that both H and $G-H$ are multi-vertex members of Γ . So, by induction on $|G|$, both H and $G-H$ are $\Delta_{\Phi(\Gamma)+1}$ -regular and hence so is G .

We may assume that G is not a member of $\Delta_{\Phi(\Gamma)+1}$ and $|G| \geq 4$. Let x be an arbitrary vertex in G . Since G is not a star shape, there exists at least one multi-vertex connected component of $G-x$.

If $G-x$ is not connected, then all multi-vertex connected components are proper subsets of $G-x$, we can set H to be any such component.

If $G-x$ is connected, let y be an adjacent vertex to x in G . We may assume that $G-y$ is also connected. If $G-\{x, y\}$ is a connected graph, we can set H to be the induced subgraph of G on $\{x, y\}$. We therefore assume that $G-\{x, y\}$ is not connected. One then checks that any connected component of $G-\{x, y\}$ is connected to both x and y . Let J be one of such connected components. We can set H to be the induced subgraph of G on $J \cup \{x\}$. ■

Remark 2.2.2. From the above theorem, $\text{Connected}_n \approx \text{Central}_n \approx \Delta_n$, $\text{Complete}_n \approx \Delta_2$, and $\text{Tree}_n \approx$

X_n for all $n \geq 3$. ■

In view of Theorem 2.2.1, we shall focus our attention on just X_n - and Δ_n -partitions, $n \geq 2$. As it turns out, each of these partitions leads to a substantially different prime factorization theory. An X_2 -partition of a graph is simply a matching since a two-vertex class corresponds to an edge in the matching and a singleton class to a vertex exposed by the matching. The classical matching theory can be recast as the network factorization theory with respect to X_2 . In such context, pertinent concepts in Chapter 1, such as, pole, root, zero, etc., become X_2 -pole, X_2 -root, X_2 -zero, etc. The next three chapters shall deal with network factorization theories with respect to Δ_2 , Δ_n and X_n , $n \geq 3$. They may be viewed as generalizations of the classical matching theory.

In the terminology of X_2 -partition, the Gallai-Edmond structure theorem becomes:

Theorem 2.2.3 (Structure theorem of X_2 -partition). *Given a graph G , let P denote the set of X_2 -poles and R the set of X_2 -roots. Then, we have*

- (a) $G-(P \cup R)$ is an X_2 -regular graph, on which every maximum matching of G induces a perfect matching.
- (b) Every connected component of the induced subgraph on P is an X_2 -blossom. Moreover, every X_2 -root is adjacent to two such X_2 -blossoms.
- (c) R is a primary X_2 -factorizer.
- (d) Let F be the induced subgraph of G on $P \cup R$. Then every vertex in P (resp. in R) is an X_2 -pole (resp. X_2 -root) of the graph F . Moreover, $\dim(G, X_2) = \dim(F, X_2)$.
- (e) $\nu(G) = (|V| - c(P) + |R|)/2$, where $c(P)$ denotes the number of odd components of the induced subgraph on P .

The Berge Formula in Chapter 1 can be restated as:

Theorem 2.2.4. For any graph G , $\dim(G, X_2) = \max\{c(G-S)-|S|: S \subset V\}$, where $c(G-S)$ denotes the number of odd components in the subgraph $G-S$.

Section 2.3. Related work on network partition

Network partition with respect to a given template is a generalization of the classical matching; meanwhile, there are many other type of generalizations. A rather relevant (to this work) approach is to generalize the matching theory by replacing the edges in a matching by some prescribed shapes; more precisely, for a given G and template Γ , a Γ -packing of G means a classification of its vertices into classes such that the induced subgraph on each class has a *spanning subgraph* (subgraph with the same vertex set with the original graph) which is isomorphic to a member of Γ . A vertex is *exposed* with respect to a Γ -packing if the vertex is a singleton class under the corresponding partition. The Γ -packing problem refers to finding a Γ -packing of a given graph that exposes the minimum number of vertices. Note that if Γ only consists of an edge, then the Γ -packing problem is precisely the classical maximum matching problem. The Γ -packing problem has been extensively studied by many authors. Comprehensive surveys in this direction can be found in [10], [25] and [34]; prominent representatives of obtained results include [2], [3], [4], [9], [11], [21], [23], [24], [26], [27], [31], [32], [33], [35], [36], [37], [48] and [49].

Recall that for a given graph G and template Γ , a Γ -partition of the graph means a partition of its vertices into classes such that the induced subgraph on each class is isomorphic to a member of Γ . In stark contrast to Γ -packing, Γ -partition has received little attention and has been investigated by only a few authors: Saito and Watanabe [47] characterized graphs with a perfect Γ -partition for the case $\Gamma = \{\text{Star}_k: k = 1, 2, \dots\}$. The case $\Gamma = X_n$ was first studied by Egawa, Kano and Kelmans [14], [29] who gave a polynomial algorithm, a Gallai–Edmonds type structure theorem and a Tutte

type theorem. In a more general setting, Király and Szabó [30] also obtained a Gallai–Edmonds type structure theorem and a Tutte type theorem, thus generalizing the work in [14], [29]. Here, we remark that these authors referred to a Γ -partition as an *induced Γ -packing*.

As a long overdue complete version of [45], this work focuses on the cases $\Gamma = X_n, \Delta_n$, to which every finite and hereditary template can be reduced (see Theorem 2.2.1). For the purpose of network partition, we note that X_n, Δ_n are in fact special cases considered in [30], so theorems obtained in [30] hold for these two families of templates as well. However, by treating each individual template with extra “care”, we obtain “finer” Gallai–Edmonds type structure theorems and Tutte type theorems for these two families of templates. A key observation in this work is that the classical Edmonds matching algorithm can be “slightly” modified (compared to algorithms proposed in [29], [30]) to adapt to the templates X_n, Δ_n for the solution to the network partition problem. More importantly, our algorithmic approach, as exemplified in Chapter 1 and further elaborated in subsequent chapters, naturally reveals ways to factorize a given graph into prime components and leads to a network factorization theory.

Chapter 3. Prime Factorization with Respect to the Template Δ_2

As described in Chapter 2, Templates Connected_n , Central_n , Complete_n , and Tree_n arise in partitioning networks of service centers via interflow connections. Theorem 2.2.1 asserts that each of these templates, as well as any other template, is equivalent to either X_n or Δ_n for some $n \geq 2$. Thus, network factorization with respect to a generic template is reduced to just an X_n - or Δ_n -partition. As it turns out, each of these partitions leads to a substantially different prime factorization theory.

Chapter 2 has recast the classical matching theory as the network factorization theory with respect to X_2 . While the next two chapters shall investigate Δ_n - and X_n -partitions, $n \geq 3$, the present chapter deals with Δ_2 -partition, which pertains to the problem of partitioning a network of service centers into groups of size up to 3 such that interflow exists between any two centers in one class. To begin with, we reiterate some basic definitions and notation about the Δ_2 -partition:

- ♦ A Δ_2 -partition of a graph G divides the vertices of G into classes such that the induced subgraph on each class is isomorphic to a singleton, an edge (that is, a pair of adjacent vertices), or a triangle K_3 .
- ♦ Given a graph G , the minimum number of vertices exposed by a Δ_2 -partition is called the Δ_2 -dimension of G , denoted by $\dim(G, \Delta_2)$.
- ♦ When $\dim(G, \Delta_2) > 0$, we say that the graph G is Δ_2 -singular. When $\dim(G, \Delta_2) = 0$, we say that G is Δ_2 -regular.
- ♦ If a particular Δ_2 -partition defines exactly $\dim(G, \Delta_2)$ singletons, that partition is called a maximum Δ_2 -partition on G . In the case of a Δ_2 -regular graph, a maximum Δ_2 -partition does not expose any vertex and is therefore called a perfect Δ_2 -partition.

- ♦ The Δ_2 -order of a vertex x in a graph G is defined as $\dim(G-x, \Delta_2) - \dim(G, \Delta_2)$. The Δ_2 -order of a vertex is always greater than or equal to -1 . A vertex with Δ_2 -order equal to -1 is called a Δ_2 -pole. A necessary and sufficient condition for a vertex to be a Δ_2 -pole is it being exposed by a maximum Δ_2 -partition.
- ♦ The maximum Δ_2 -order of any vertex is $\Phi(\Delta_2) = 1$. A vertex with Δ_2 -order equal to 1 is called a Δ_2 -zero.
- ♦ If a vertex is not a Δ_2 -pole vertex but is adjacent to at least one Δ_2 -pole, then it is called a Δ_2 -root.
- ♦ A set S of vertices is a Δ_2 -factorizer of G if there are $\dim(G, \Delta_2) + \Phi(\Delta_2)|S|$ connected components of $G-S$ that are Δ_2 -singular graphs. Since $\Phi(\Delta_2) = 1$, this implies that exactly $\dim(G, \Delta_2) + |S|$ connected components of $G-S$ have Δ_2 -dimension 1 and all others are Δ_2 -regular graphs.
- ♦ A Δ_2 -prime graph is defined as a connected graph that has no non-empty Δ_2 -factorizer.
- ♦ A Δ_2 -singular Δ_2 -prime graph is called a Δ_2 -blossom.
- ♦ A Δ_2 -factorizer S is called a primary Δ_2 -factorizer if all Δ_2 -singular components of $G-S$ are Δ_2 -blossoms (while the Δ_2 -regular components may or may not be Δ_2 -prime).
- ♦ A Δ_2 -factorizer S of a graph is called a prime Δ_2 -factorizer if all connected components of $G-S$ are Δ_2 -prime graphs.

Section 3.1. An Edmonds-type algorithm

Definition 3.1.1. A path in a graph is said to be an *alternating path* with respect to a Δ_2 -partition M if pairs of adjacent vertices on the path are alternately classmates and non-classmates under M . An alternating path (x_0, x_1, \dots, x_k) with respect to a Δ_2 -partition M is called an *augmenting path* with

respect to M , if the following conditions are satisfied:

1. x_0 is exposed by M .
2. If k is an odd integer, then the class of x_k defined by M is either a singleton or a triangle.
3. If k is an even integer, then for some $m < k/2$, both x_{2m} and x_{2m+1} are adjacent to x_k .




Figure 3-1: This figure displays for types of Δ_2 -augmenting path (x_0, x_1, \dots, x_k) with respect to a Δ_2 -partition M . In (a) and (b), the path length $k = 5$ is odd and the class of x_k defined by M is either a singleton or a triangle, respectively. In (c) and (d), the path length $k = 6$ is even and both x_{2m} and x_{2m+1} are adjacent to x_k , where $m = k/2 - 1$ and $m < k/2 - 1$, respectively.

The following algorithm is the Δ_2 -partition counterpart of Algorithm 1.2.6.

Algorithm 3.1.2. Given a Δ_2 -partition M on a graph G , this algorithm determines whether G admits an augmenting path with respect to M . Write $G_0 = G$ and $M_0 = M$. The algorithm will construct a sequence of graphs G_t , $0 \leq t \leq \tau$, with a Δ_2 -partition M_t on each G_t . In the end, whether there is an

augmenting path with respect to M_t in G_t will be apparent. If there is, then, for every t , an augmenting path with respect to M_{t+1} in G_{t+1} induces an augmenting path with respect to M_t in G_t . The graph G_t will be associated with, besides the matching M_t , an acyclic subgraph T_t , in which every vertex is labeled either *even* or *odd* so that T_t is a bipartite graph between even and odd vertices. Figure 3-2 illustrates G_t , M_t and T_t for a generic t .

Initially, those vertices z_1, z_2, \dots, z_d exposed by M are all labeled as even vertices. Let T_0 consist of these d vertices. Given G_t , M_t and T_t , the corresponding iterative step in the algorithm achieves exactly one of the following:

- (a) Keep both G_t and M_t the same, whereas grow T_t by adding an odd vertex, an even vertex, and two edges. The first edge is between an existing even vertex and the new odd vertex; the second is between the new vertices and belongs to M_t . At the end of this step, increase the index t by 1.
- (b) Contract an odd cycle in T_t (and G_t) to obtain T_{t+1} (and G_{t+1}), and let M_t induce a Δ_2 -partition M_{t+1} on G_{t+1} . At the end of this step, increase the index t by 1.
- (c) Identify an augmenting path of M_t , and recursively find an augmenting path with respect to M . The algorithm terminates, that is, t is the final index τ .
- (d) G_t does not admit any augmenting path with respect to M_t , and hence G does not admit any augmenting path with respect to M . The algorithm terminates.

The iterative step at time t starts by looking for an edge of G_t such that it is

- not an edge of T_t ,
- incident to at least one even vertex of T_t , and
- is not incident to any odd vertex of T_t .

We then have the following cases:

Case 1. Such an edge does not exist. Then G does not admit an augmenting path with respect to M . The algorithm terminates. ((d) is achieved.)

Case 2. Such an edge exists. Let (e, f) be such an edge of G_t , where e is an even vertex of G_t .

Case 2.1. f is not in T_t .

Case 2.1.1. The class of f defined by M_t consists of three vertices. It is easy to see that there is an odd-length augmenting path in G_t with respect to M_t , thus implying the existence of an augmenting path in G with respect to M by Lemma 3.1.3. Then the algorithm terminates; see Figure 3-3 ((d) is achieved).

Case 2.1.2. The class of f defined by M_t consists of exactly two vertices, say f and g (necessarily g is outside T_t). In this case, we add two vertices f and g , and two edges (e, f) and (f, g) in T_t to obtain T_{t+1} . The vertex f is labeled odd and g even. Set $G_{t+1} = G_t$, $M_{t+1} = M_t$. Increase t by 1; see Figure 3-4 ((a) is achieved).

Case 2.2. f is an even vertex in T_t . Then there exists a unique path $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n} = e)$ in T_t connecting e to an exposed vertex by M_t , where (x_{2i-1}, x_{2i}) forms a class of M_t for $1 \leq i \leq n$. Similarly, there exists a path $(y_0, y_1, y_2, \dots, y_{2m-1}, y_{2m} = e)$ such that y_0 is an exposed vertex by M_t and $\{y_{2j-1}, y_{2j}\}$ forms a class of M_t for $1 \leq j \leq m$. We further consider the following two subcases.

Case 2.2.1. $x_0 = y_0$. Let $k \geq 0$ be the largest index with $x_k = y_k$ (k must be an even integer). Thus, $(x_k, x_{k+1}, \dots, x_{2n} = e, y_{2m} = f, \dots, y_{k+1}, y_k = x_k)$ is an odd cycle, which corresponds to a blossom, say B , in G . We consider two subcases depending on whether this blossom is Δ_2 -singular or not.

Case 2.2.1.1. B is Δ_2 -singular. Contract the cycle $(x_k, x_{k+1}, \dots, x_{2n} = e, y_{2m} = f, \dots, y_{k+1})$ into a single vertex to obtain G_{t+1} , and set M_{t+1} to be the induced partition by M_t on G_{t+1} . Increase t by 1; see Figure 3-5 ((b) is achieved).

Case 2.2.1.2. B is Δ_2 -regular. We then construct an augmenting path in G with respect to M and thereby terminate the algorithm. For this case, M induces a maximum X_2 -partition, say O , on B . One then checks (see Theorem 3.2.14) that there exists a perfect Δ_2 -partition, say Q' , on B with exactly one K_3 -class, say $\{u, v, w\}$. Let Q be the X_2 -partition replacing the K_3 -class of Q' with three singleton classes. Apparently, there is only one vertex, say s , exposed by O , while there are three by Q . Applying Lemma 1.2.1, we deduce the existence of a path of type 2 and another of type 4. These paths are alternating paths with respect to both O and Q . One of the two paths, say P_1 , is of even length, connecting s to an exposed vertex, say u , by Q ; the other path, say P_2 , is of odd length, connecting the remaining two exposed vertices by Q , say v and w . Moreover, the beginning and ending edges of P_2 both form the same classes of O . Also, there exists an alternating path in G , say P_3 , of even length, with respect to M , connecting an exposed vertex by M to s . Gluing the paths P_3 , P_1 , (u, v) and P_2 together gives us an even-length augmenting path (satisfying condition 3 in Definition 3.1.1) in G with respect to M ; see Figure 3-6 ((d) is achieved).

Case 2.2.2. $x_0 \neq y_0$. Then $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n} = e, y_{2m} = f, y_{2m-1}, \dots, y_1, y_0)$ is an augmenting path in G_t with respect to M_t , which implies an augmenting path with respect to M in G by Lemma 3.1.3. The algorithm terminates; see Figure 3-7 ((c) is achieved). ■

Lemma 3.1.3. *For any t , there is an augmenting path in G_t with respect to M_t if and only if there is an augmenting path in G_{t+1} with respect to M_{t+1} .*

Proof. The proof is similar to that of Lemma 1.2.7, thus omitted. ■




Figure 3-2: G_t , T_t , and M_t are constructed in Algorithm 3.1.2 by time t . An even vertex of T_t is represented by a rectangle, an odd vertex of T_t by a hollow circle, and a vertex in $G_t - V(T_t)$ by a solid circle. An edge of G_t is regarded as outside T_t if it is incident with a vertex outside T_t . Classes in M_t are indicated by highlighted edges. The figure also displays (inside rectangles) those multi-vertex blossoms in G that have been contracted into even vertices of T_t .




Figure 3-3: Illustration for Case 2.1.1




Figure 3-4: Illustration for Case 2.1.2




Figure 3-5: Illustration for Case 2.2.1.1




Figure 3-6: Illustration for Case 2.2.1.2, where the augmenting path is highlighted as a dotted path




Figure 3-7: Illustration for Case 2.2.2, where the augmenting path is highlighted as a dotted path

Section 3.2. Prime factorization of networks with respect to Δ_2

The next theorem is a Δ_2 -partition counterpart of Theorem 1.3.11.

Theorem 3.2.1 (Structure theorem of Δ_2 -partition). *For a graph G , let P denote the set of Δ_2 -poles, and R the set of Δ_2 -roots. Then we have*

- (a) $G - (P \cup R)$ is a Δ_2 -regular graph;
- (b) Every connected component of the induced subgraph on P is a Δ_2 -singular blossom.

Moreover, every Δ_2 -root is adjacent to at least two such Δ_2 -singular blossoms;

- (c) R is a primary Δ_2 -factorizer;

(d) Let F be the induced subgraph of G on $P \cup R$. Then every vertex from P (resp. from R) is a Δ_2 -pole (resp. Δ_2 -root) of the graph F . Moreover, $\dim(G, \Delta_2) = \dim(F, \Delta_2)$.

Proof. The proof is trivial when G is Δ_2 -regular, so we only consider the case when G is Δ_2 -singular. Consider any Δ_2 -partition M on G such that there is no augmenting path in G with respect to M , and let z_1, z_2, \dots, z_d denote the vertices exposed by M . Apply Algorithm 3.1.2 on G with respect to M . It can be easily checked that, at any time t , the following 5 basic properties are satisfied:

- (1) Every odd vertex in T_t is a vertex of the original graph G , and so is every vertex in $G_t - V(T_t)$.
Every even vertex in T_t corresponds to a contracted Δ_2 -singular blossom in G .
- (2) If $\{f, g\}$ is a two-vertex class defined by M_t , then either both f and g or neither of them are vertices in $G_t - V(T_t)$. Moreover, vertices in any three-vertex class defined by M_t are outside T_t .
- (3) Each odd vertex in T_t is adjacent to exactly two even vertices and belongs to a two-vertex class defined by M_t .
- (4) Each connected component of T_t contains exactly one exposed vertex by M_t .
- (5) The number of even vertices in T_t exceeds the number of odd vertices exactly by d .

We deduce from the above five properties the sixth property:

- (6) In the original graph G , there exists an odd-length (resp. even-length) alternating path, with respect to M , from an exposed vertex by M to an odd vertex in T_t (resp. a vertex in a blossom in G corresponding to an even vertex of T_t).

We shall only prove the “odd vertex” part of (6), the other part being similar. Let x be an odd

vertex in T_t . From Property (4), there exists a unique odd-length alternating path in T_t with respect to M_t from an exposed vertex by M_t to x . Let this path be $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n}, x_{2n+1} = x)$. By Property (3), the pair $\{x_{2i-1}, x_{2i}\}$ forms a class of M_t for $1 \leq i \leq n$. For $0 \leq i \leq n$, let B_{2i} be the blossom in G corresponding to the even vertex x_{2i} and s_{2i} the exposed vertex by M on B_{2i} and t_{2i} a vertex in B_{2i} that is adjacent to the odd vertex x_{2i+1} . From Lemma 1.3.9, there exists an even-length alternating path in G with respect to M from s_{2i} to t_{2i} , for $0 \leq i \leq n$. These alternating paths and the paths $(t_{2i}, x_{2i+1}, s_{2i+2})$, $0 \leq i \leq n$, and (t_{2n}, x_{2n+1}) can be pieced together to form an odd-length alternating path in G with respect to M from the exposed vertex s_0 by M to the odd vertex x_{2n+1} .

Since T_t is bipartite, we conclude that

(7) In G_t , every even vertex is adjacent to only odd vertices.

Since G does not admit any augmenting path with respect to M , Algorithm 3.1.2 can only terminate in Case 1.

Let R'_τ denote the set of odd vertices in T_τ . From Property (5), there are $d+|R'_\tau|$ even vertices. And by Property (7), every even vertex is by itself a connected component in $G_\tau - R'_\tau$. Thus, $\dim(G_\tau, \Delta_2) \geq d$. On the other hand, M_τ exposes exactly d vertices in T_τ and none in $G_\tau - V(T_\tau)$. We therefore reach the following conclusions:

(8) $\dim(G_\tau, \Delta_2) \geq d$.

(9) M_τ is a maximum Δ_2 -partition on G_τ .

Now, consider any even vertex e in T_τ . From Properties (5) and (3), there exists in G_τ an alternating path with respect to M_τ from an exposed vertex by M_τ to e . This proves that e is a Δ_2 -pole of G_τ . We thus proved:

(10) Every even vertex is a Δ_2 -pole of G_τ .

From Property (1), vertices in R'_τ are also vertices of the original graph G and there are at least $d+|R'_\tau|$ components in $G-R'_\tau$ that are Δ_2 -singular blossoms. Thus, $\dim(G, \Delta_2) \geq d$. By the same argument, if x is a vertex in $G_\tau-V(T_\tau)$ (and hence also a vertex of G , by Property (1)), then $\dim(G-x, \Delta_2) \geq d$. On the other hand, the Δ_2 -partition M on G exposes exactly d vertices. We therefore reach the following conclusions:

$$(11) \quad \dim(G, \Delta_2) = d.$$

$$(12) \quad \text{If } x \text{ is a vertex in } G_\tau-V(T_\tau), \text{ then } x \text{ is not a } \Delta_2\text{-pole of } G.$$

$$(13) \quad R'_\tau \text{ is a } \Delta_2\text{-factorizer of } G.$$

$$(14) \quad M \text{ is a maximum } \Delta_2\text{-partition on } G.$$

We next prove that

$$(15) \quad \text{A vertex of } G \text{ is a } \Delta_2\text{-pole if and only if it belongs to a } \Delta_2\text{-singular blossom in } G \text{ that is contracted into an even vertex of } T_\tau.$$

From Property (13) and Lemma 2.1.3, every odd vertex of T_τ is a Δ_2 -zero of G . This, together with Property (12), proves the “only if” part of Property (15). Conversely, let B be a Δ_2 -singular blossom in G that is contracted into an even vertex e' of T_τ and let e be a vertex in B . We need to show that e is a Δ_2 -pole of G . From Properties (8) and (10), we know that $\dim(G_\tau-e', \Delta_2) = d-1$. Note that the Δ_2 -dimension of a graph is unchanged when a Δ_2 -singular blossom in it is contracted to a single vertex such that this vertex is a Δ_2 -pole in the contracted graph. Therefore $\dim(G-B, \Delta_2) = \dim(G_\tau-e', \Delta_2)$. Thus

$$\begin{aligned} \dim(G-e, \Delta_2) &\leq \dim(G-B, \Delta_2) + \dim(B-e, \Delta_2) = d-1+0 \\ &= \dim(G, \Delta_2) - 1, \quad \text{by Property (11)}. \end{aligned}$$

Property (15) is then proved.

By Properties (13), (1) and Lemma 2.1.3, every odd vertex is a Δ_2 -zero of G . Thus every vertex in R'_τ is a Δ_2 -root of G , by Properties (3) and (15). On the other hand, by Properties (7) and (15), there exist no Δ_2 -roots other than the odd vertices. Therefore, we have proved that

(16) A vertex of G is a Δ_2 -root if and only if it is an odd vertex of T_τ , that is, $R'_\tau = R$.

(17) The induced subgraph of G on the vertices of $G_\tau - V(T_\tau)$ is Δ_2 -regular graph.

Statement (a) of Theorem 3.2.1 is due to Properties (15), (16) and (17). Statement (b) is due to Properties (15) and (3). From Properties (13) and (16), we know that R is a Δ_2 -factorizer of G . Furthermore, by Properties (7), (1) and (17), every Δ_2 -singular connected component of $G - R$ is a Δ_2 -singular blossom, which is a Δ_2 -prime graph, by Corollary 3.2.11. Thus (c) is proved.

Now we turn to prove (d).

From (a), the graph $G - P \cup R$ is Δ_2 -regular. Thus,

$$d = \dim(G, \Delta_2) \leq \dim(G - P \cup R, \Delta_2) + \dim(F, \Delta_2) = \dim(F, \Delta_2).$$

On the other hand, the Δ_2 -partition M induces a Δ_2 -partition on F which exposes d vertices, that is, $\dim(F, \Delta_2) \leq d$. Therefore, we have

$$(18) \quad \dim(F, \Delta_2) = \dim(G, \Delta_2) = d.$$

Let x be a vertex in P , that is, $\dim(G - x, \Delta_2) = d - 1$. Consider a maximum Δ_2 -partition M^* on G which exposes x . From Properties (15) and (16), there exists no Δ_2 -pole of G in $G - P \cup R$, thus the M^* induces a Δ_2 -partition on F which also exposes x and some other $d - 1$ vertices. Thus

$$\dim(F - x, \Delta_2) \leq \dim(G - x, \Delta_2) = d - 1 = \dim(F, \Delta_2) - 1,$$

which implies that x is also a Δ_2 -pole of F , that is,

(19) Every vertex in P is a Δ_2 -pole of the graph F .

From Statements (a) and (c), R is also a Δ_2 -factorizer of F . Thus every vertex in R is a Δ_2 -zero of F . Therefore

(20) Every vertex in R is a Δ_2 -root of the graph F .

Statement (d) is then established. ■

Property (14) in the above proof yields the following theorem as a byproduct.

Theorem 3.2.2. *With respect to any non-maximum Δ_2 -partition on a given graph, there exists at least one Δ_2 -augmenting path.*

The converse of Theorem 3.2.2 is obviously true:

Lemma 3.2.3. *If (x_0, x_1, \dots, x_k) is a Δ_2 -augmenting path with respect to a Δ_2 -partition M , then there exists another Δ_2 -partition that covers x_0 and all vertices covered by M .*

Example 3.2.4. The Δ_2 -partition of a graph G in Figure 3-8 is not maximum. In fact, the vertex sequences $(7, 8, 11, 12, 14, 16)$ and $(16, 14, 12, 13, 15, 18)$ are two Δ_2 -augmenting paths of odd length, while the vertex sequence $(10, 11, 8, 6, 5, 3, 1, 2, 4)$ is a Δ_2 -augmenting path of even length.

■




Figure 3-8: A non-maximum Δ_2 -partition on a graph.

A graph $B = (V, E)$ is called a Δ_2 -bud if it admits no perfect Δ_2 -partition but, for some $v \in V$, $B - v$ admits a perfect matching. In particular, a singleton graph is also a Δ_2 -bud. The following theorem is a Δ_2 -partition counterpart of Theorem 1.3.13, the Berge formula for the matching theory.

Theorem 3.2.5. *For a vertex subset S of a graph G , let $b(G-S)$ denote the number of components of $G-S$ which are Δ_2 -buds. Then, $\dim(G, \Delta_2) = \max_{S \subseteq V} \{b(G-S) - |S|\}$.*

Proof. We first prove that $\dim(G, \Delta_2) \geq \max_{S \subseteq V} \{b(G-S) - |S|\}$. Let M be a maximum Δ_2 -partition of G . Let S be any vertex subset, and let G_1, \dots, G_k , $k := b(G-S)$, denote all the Δ_2 -bud components of $G-S$. Among these components, renumbering if necessary, let G_1, \dots, G_j be those containing a vertex exposed by M . Then for each $j+1 \leq i \leq k$, there is at least one edge in M from S to G_i , which implies $|S| \geq k - j$. Apparently, $\dim(G, \Delta_2) \geq j$ since each of G_1, \dots, G_j contains an exposed vertex. Hence $\dim(G, \Delta_2) \geq j \geq k - |S| = b(G-S) - |S|$. We then conclude that $\dim(G, \Delta_2) \geq \max_{S \subseteq V} \{b(G-S) - |S|\}$. On the other hand, apply Algorithm 3.1.2 to G , by Property (5) in the proof of Theorem 3.2.1, the even vertices in T_τ outnumber the odd vertices exactly by d . Since an even vertex in T_τ corresponds to a Δ_2 -singular blossom (which is a Δ_2 -bud) in G , it then follows that if we choose S to be R , the set of Δ_2 -roots of G , we have $b(G-R) - |R| = \dim(G, \Delta_2)$. This establishes $\dim(G, \Delta_2) = \max_{S \subseteq V} \{b(G-S) - |S|\}$. ■

The following theorem is an immediate corollary of Theorem 3.2.5, which gives a necessary and sufficient condition on the existence of a perfect Δ_2 -partition.

Theorem 3.2.6. *A graph G admits a perfect Δ_2 -partition if and only if for every vertex subset S , $b(G-S) \leq |S|$ holds.*

The following theorem is the Δ_2 -counterpart to Mendelsohn-Dulmage Theorem. It follows

directly from Theorem 3.2.1 and Lemma 3.2.3.

Theorem 3.2.7. *For any Δ_2 -partition M on a graph G , there exists a maximum Δ_2 -partition covering all the vertices of G covered by M .*

Remark 3.2.8. Let B be a graph with a vertex v such that $B-v$ has a perfect matching. If there is a triangle subgraph K such that $B-K$ has a perfect matching, then clearly B admits a perfect Δ_2 -partition. The opposite is also true as explained below. The perfect matching on $B-v$ is a non-perfect Δ_2 -partition on B with the unique exposed vertex v . If B admits a perfect Δ_2 -partition, then, according to Theorem 3.2.2, there exists a Δ_2 -augmenting path (satisfying Condition 2 or 3 in Definition 3.1.1) with respect to the said non-perfect Δ_2 -partition. This augmenting path then leads to a triangle subgraph K such that $B-K$ has a perfect matching. ■

Remark 3.2.9. Let \mathcal{G} be a family of connected graphs. A \mathcal{G} -packing of a graph G is a subgraph S of G such that each connected component of S is isomorphic to a member of \mathcal{G} . \mathcal{G} -packing theory is proposed by P. Hell and D. G. Kirkpatrick in 1981 (see [24], [25],[26],[27] and [31]) as a natural generalization of the classical matching theory. Recall that a Γ -partition of a graph $G = (V, E)$ divides the vertex set V into subclasses V_1, V_2, \dots, V_k such that the induced subgraph of G on each V_i is isomorphic to a graph from the family Γ . While on the other hand, a \mathcal{G} -packing defines a set of disjoint subgraphs (not required to be induced) such that each such subgraph is isomorphic to a member graph of the family. Since an induced subgraph $D = (V(D), E(D))$ may have more edges than a subgraph $D' = (V(D), A(D'))$ with the same vertex set $V(D)$, a Γ -packing may not be a Γ -partition. However, apparently, a Γ -partition defines a Γ -packing; in particular, a Δ_2 -partition is just a $\{K_2, K_3\}$ -packing.

We next state results on Δ_2 -prime graphs. We start with blossoms (Definition 1.3.7).

Lemma 3.2.10. *A blossom is a Γ -prime graph for any template $\Gamma \supset X_2$.*

Proof. Let G be a blossom and x be an arbitrary vertex in G . From Lemma 1.3.8, $\dim(G-x, X_2) = 0$. It then follows from $\Gamma \supset X_2$ that $\dim(G-x, \Gamma) \leq \dim(G-x, X_2) = 0$. Thus x is not a Γ -zero of G . By Lemma 2.1.3, G is a Γ -prime graph. ■

Corollary 3.2.11. *A blossom is a Δ_2 -prime graph.*

It immediately follows from Lemma 3.2.10 and Lemma 1.3.8 that

Corollary 3.2.12. *Every vertex in a Δ_2 -regular blossom has Δ_2 -order 0, and every vertex in a Δ_2 -singular blossom is a Δ_2 -pole.*

The converse of Corollary 3.2.12 is not true in general but will later be established for Δ_2 -singular graphs. For this reason, we are interested in the characterization of Δ_2 -singular blossoms, or equivalently, the characterization of Δ_2 -regular blossoms.

Lemma 3.2.13. *If G is a Δ_2 -regular blossom, then there exists a perfect Δ_2 -partition on G that defines exactly one K_3 -class such that, when this K_3 is contracted into a single vertex, the graph G becomes another blossom.*

Proof. Assume that, among all perfect Δ_2 -partitions on G , the Δ_2 -partition P defines the minimum number of K_3 -classes. This minimum number must be odd since the order of a blossom is odd and every other class defined by P consists of two adjacent vertices. Let G' denote the graph resulting from G by contracting every K_3 -class defined by P into a single vertex. Thus P induces an X_2 -partition P' on the graph G' . The K_3 -classes of P are in a one-to-one correspondence with the exposed vertices defined by P' . Note that G' does not admit any X_2 -augmenting path with respect to P' because otherwise there would exist a perfect Δ_2 -partition on G which defines two fewer

K_3 -classes than P . It then follows that P' is a maximum X_2 -partition on the graph G' . Hence every exposed vertex of P' is an X_2 -pole of G' .

We claim that if S is an X_2 -factorizer of G' then it is also an X_2 -factorizer of G . In fact, by Lemma 2.1.3, S is a set of X_2 -zeros thus S does not contain any exposed vertex of P' . Consequently, a connected component of odd order in the graph $G'-S$ is also that of $G-S$ and vice versa. Thus, S is also an X_2 -factorizer of G .

It now remains to show that P defines only one K_3 -class and that G' is a blossom. Since G is a blossom, G contains no non-empty X_2 -factorizer. Then, by the above claim, G' does not contain any non-empty X_2 -factorizer, and thus is X_2 -prime. It then follows from Theorem 1.3.15 that G' is also a blossom. The maximum X_2 -partition P' of the blossom G' defines only one exposed vertex, thus P defines only one K_3 -class. ■

Theorem 3.2.14. *A blossom G is Δ_2 -regular if and only if there exists a K_3 -subgraph whose contraction into a single vertex would transform G into another blossom.*

Proof. The “only if ” part follows from Lemma 3.2.13. For the “if ” part, we assume the contraction of a K_3 -subgraph H into a single vertex turns G into a smaller blossom. From Lemma 1.3.8, a blossom has X_2 -dimension equal to 1 and every vertex in the blossom is an X_2 -pole. Thus the graph $G-H$ is X_2 -regular and hence also Δ_2 -regular. Thus G is Δ_2 -regular too. ■

It should be pointed out that a blossom containing exactly one K_3 -subgraph may be Δ_2 -singular. Such an example is shown in Figure 3-9.




Figure 3-9: A Δ_2 -singular blossom.

The following theorem characterizes all Δ_2 -singular Δ_2 -prime graphs.

Theorem 3.2.15 (Characterization of Δ_2 -blossom). *The following statements are equivalent for a graph G .*

- (a) G is a Δ_2 -blossom, that is, a Δ_2 -singular Δ_2 -prime graph.
- (b) G is a Δ_2 -prime graph with Δ_2 -dimension 1.
- (c) G is connected and all vertices are Δ_2 -poles.
- (d) G is a Δ_2 -singular blossom.

Proof. (a) \Rightarrow (b). Assume for contradiction that G has Δ_2 -dimension greater than 1. Let M be a maximum Δ_2 -partition of G with exposed vertices z_1, z_2, \dots, z_d , where $d > 1$. Apply Algorithm 3.1.2

to G with respect to M . Then, there are odd vertices in T_τ , and all the odd vertices constitute a non-empty factorizer by Statement (c) of Theorem 3.2.1, a contradiction.

(b) \Rightarrow (c). Let M be a maximum Δ_2 -partition of G with exposed vertices z_1 . Apply Algorithm 3.1.2 to G with respect to M . Then there are no odd vertices in T_τ , thus T_τ consists of only one even vertex, which implies that all vertices are poles.

(c) \Rightarrow (d). By Theorem 3.2.1, all poles, thus all vertices, in G are contained in blossoms, which will be contracted into even vertices when Algorithm 3.1.2 terminates. Since there are no roots in G , all the vertices constitute exactly one blossom.

(d) \Rightarrow (a). This follows from Corollary 3.2.11. ■

Remark 3.2.16. It follows from Theorem 3.2.15 that Δ_2 -singular blossoms are the only Δ_2 -prime graphs that are Δ_2 -singular. Meanwhile, there are different kinds of Δ_2 -prime graphs that are Δ_2 -regular. For example, any complete graph of order larger than 2 is a Δ_2 -regular Δ_2 -prime graph. ■

Theorem 3.2.17. *For any graph G , every primary Δ_2 -factorizer contains all Δ_2 -roots. Thus, the set of Δ_2 -roots is the unique minimal primary Δ_2 -factorizer.*

Proof. Let M be a maximum Δ_2 -partition of G , isolating vertices z_1, z_2, \dots, z_d . Apply Algorithm 3.1.2 to G with respect to M . Then, R , the set of Δ_2 -roots of G , is just the set of odd vertices in T_τ . Suppose, by contradiction, that there is a primary Δ_2 -factorizer S such that S does not contain R as a subset. Pick an odd vertex $x \in R \setminus S$. One then checks that the component C of $G_\tau - S$ containing x has more even vertices than odd vertices. Since even vertices in G_τ correspond to Δ_2 -singular blossoms in the original graph G , C must be a singular component of $G - S$. But, by Theorem 3.2.15, C is not prime, which contradicts the assumption that S is a primary factorizer. ■





Figure 3-10: An example graph

Example 3.2.18. Let G be the graph in Figure 3-10. It can be checked that the Δ_2 -pole set $P = \{1, 2, 3, 4, 5, 7, 8\}$, Δ_2 -root set $R = \{6\}$ and vertex 6 is the only Δ_2 -zero. The induced subgraph on $G - (P \cup R)$ is a triangle consisting of the vertices $\{9, 10, 11\}$, which is clearly Δ_2 -regular. The connected components of the induced subgraph on P are either the single vertex sets $\{7\}$ and $\{8\}$ or the cycle of order 5, which are of course Δ_2 -singular blossoms. The set $\{6\}$ is the unique primary Δ_2 -factorizers of G . Since the triangle $\{9, 10, 11\}$ is also a Δ_2 -prime graph, $\{6\}$ is also a prime Δ_2 -factorizer of G . The induced subgraph, say F , of G on $P \cup R$ is shown in Figure 3-10. It is easy to see that $\dim(G, \Delta_2) = \dim(F, \Delta_2) = 2$ and vertex 6 is the only Δ_2 -root of F and the other vertices in F are all Δ_2 -poles. ■

Figure 3-11: The induced subgraph F of G on $P \cup R$.

Chapter 4. Prime Factorization with Respect to the Template Δ_n

This chapter concentrates on network factorization theory with respect to Δ_n , here and throughout this chapter, n stands for an arbitrary yet fixed integer greater than or equal to 3.

We first reiterate state some basic definitions and notation on Δ_n -partition.

- ◆ A Δ_n -partition of a graph G divides the vertices of G into classes such that the induced subgraph on each class is isomorphic to a triangle K_3 or a star-shaped graph $Star_k$, $1 \leq k \leq n$.
- ◆ For a graph G , the minimum number of vertices exposed by a Δ_n -partition is called the Δ_n -dimension of G , denoted by $\dim(G, \Delta_n)$.
- ◆ When $\dim(G, \Delta_n) > 0$, we say that G is Δ_n -singular. When $\dim(G, \Delta_n) = 0$, we say that G is Δ_n -regular.
- ◆ If a particular Δ_n -partition defines exactly $\dim(G, \Delta_n)$ singletons, that partition is called a maximum Δ_n -partition on G . In the case of a Δ_n -regular graph, a maximum Δ_n -partition does not expose any vertex and is therefore called a perfect Δ_n -partition.
- ◆ The Δ_n -order of a vertex x in a graph G is defined as $\dim(G-x, \Delta_n) - \dim(G, \Delta_n)$. The Δ_n -order of a vertex is always greater than or equal to -1 . A vertex with Δ_n -order equal to -1 is called a Δ_n -pole. A necessary and sufficient condition for a vertex to be a Δ_n -pole is it being exposed by a maximum Δ_n -partition.
- ◆ The maximum Δ_n -order of any vertex is $\Phi(\Delta_n) = n-1$. A vertex with Δ_n -order equal to $\Phi(\Delta_n) = n-1$ is called a Δ_n -zero.
- ◆ If a vertex is not a Δ_n -pole vertex but adjacent to at least one Δ_n -pole, then it is called a Δ_n -root.
- ◆ A vertex is called a Δ_n -infinity if all adjacent vertices (if any) are Δ_n -zeroes.
- ◆ A set S of vertices is called a Δ_n -factorizer of G if there are $\dim(G, \Delta_n) + \Phi(\Delta_n)|S|$ connected

components of $G-S$ that are Δ_n -singular graphs. Since $\Phi(\Delta_n) = n-1$, this implies that exactly $\dim(G, \Delta_n) + (n-1)|S|$ connected components of $G-S$ have Δ_n -dimension 1 and all others are Δ_n -regular graphs.

- ◆ A Δ_n -prime graph is defined as a connected graph that has no non-empty Δ_n -factorizer.
- ◆ A Δ_n -singular Δ_n -prime graph is called a Δ_n -blossom.
- ◆ A Δ_n -factorizer S is called a primary Δ_n -factorizer if all Δ_n -singular components of $G-S$ are Δ_n -blossoms (while the Δ_n -regular components may or may not be Δ_n -prime).
- ◆ A Δ_n -factorizer S of a graph is called a prime Δ_n -factorizer if all connected components of $G-S$ are Δ_n -prime graphs.

Section 4.1. An Edmonds-type algorithm

Let M be a Δ_n -partition of a graph G . A vertex x in G is said to be a *Star_n-center* with respect to M if the induced graph on the class (defined by M) of x is a $Star_n$ with x as its center. A path on a graph is said to be an *alternating path* with respect to a Δ_n -partition if pairs of adjacent vertices on the path are alternately classmates and non-classmates under the Δ_n -partition. A Δ_n -alternating path $(x_0, x_1, \dots, x_{2k+1})$ with respect to M is called a *Δ_n -augmenting path* with respect to M if the following conditions are satisfied:

1. x_0 is exposed by M .
2. The vertex x_{2k+1} is not a $Star_n$ -center with respect to M .

Lemma 4.1.1. *Let $(x_0, x_1, \dots, x_{2k+1})$ be a Δ_n -augmenting path with respect to a Δ_n -partition M . Then there exists a Δ_n -partition that covers x_0 and all vertices covered by M .*

Proof. It suffices to prove that the lemma holds for the case when $x_1, x_3, \dots, x_{2k-1}$ are all $Star_n$ -center with respect to M but x_{2k+1} is not. If x_{2k+1} is the center vertex in its class, then the class size is strictly smaller than n . In this case, we modify M to form a new Δ_n -partition by deleting x_{2k}

from its class and adding vertex x_{2k} into the class of x_{2k+1} . If x_{2k+1} is not the center vertex in its class, we modify M to form a new Δ_n -partition by deleting x_{2k}, x_{2k+1} from their classes respectively and creating a new class $\{x_{2k}, x_{2k+1}\}$. For either case, under the new Δ_n -partition, $(x_0, x_1, \dots, x_{2k-1})$ is a shorter Δ_n -augmenting path, and $x_1, x_3, \dots, x_{2k-3}$ are all Star_n -center but x_{2k-1} is not. The lemma then follows from an inductive argument. ■

Algorithm 4.1.2. Given a Δ_n -partition M on a graph G , this algorithm determines whether G admits a Δ_n -augmenting path with respect to M . The algorithm will construct a sequence of acyclic graphs T_t , $0 \leq t \leq \tau$, in which every vertex is labeled either *even* or *odd* so that T_t is a bipartite graph between even and odd vertices. Figure 3-2 illustrates T_t constructed by time t .

Initially, those vertices z_1, z_2, \dots, z_d exposed by M are all labeled as even vertices. Let T_0 consist of these d vertices. Given T_t , the corresponding iterative step in the algorithm achieves exactly one of the following:

- (a) Grow T_t by adding an odd vertex, $(n-1)$ even vertices, and n edges. The first edge is between an existing even vertex and the new odd vertex; the remaining $(n-1)$ edges are between the new odd vertex and $(n-1)$ new even vertices and belong to M . At the end of this step, increase the index t by 1.
- (b) Identify a Δ_n -augmenting path with respect to M . The algorithm terminates, that is, t is the final index τ .
- (c) G does not admit any augmenting path with respect to M . The algorithm terminates.

The iterative step at time t starts by looking for an edge of G such that it is

- not an edge of T_t ,
- incident to at least one even vertex of T_t , and
- is not incident to any odd vertex of T_t .

We then have the following cases:

Case 1. Such an edge exists, say (e, f) , where e is an even vertex of T_t and f is either an even vertex of T_t or outside T_t .

Case 1.1. The vertex f is an even vertex of T_t , or f is outside T_t but not a Star_n -center with respect to M . One then verifies that there exists a Δ_n -alternating path with respect to M from e to an exposed vertex by M , and this Δ_n -alternating path and vertex f form a Δ_n -augmenting path in G with respect to M . Then the algorithm terminates, see Figure 4-1 and Figure 4-2. ((b) is achieved.)

Case 1.2. Vertex f is outside T which is a Star_n -center with respect to M , then T_t is enlarged to T_{t+1} by adding the induced subgraph of G on the whole class of f defined by M . Among the n new vertices in T_{t+1} , the vertex f is labeled as odd and all the other $n-1$ vertices are labeled as even. Increase t by 1; see Figure 4-3 and Figure 4-4. ((a) is achieved.)

Case 2. Every edge of G that is incident with at least one even vertex but not any odd vertex is an edge of T_t . In this case, G does not admit any augmenting path; the algorithm terminates. ((c) is achieved.) ■




Figure 4-1: Illustration for Case 1.1 for $n = 4$ and f is an even vertex of T_t




Figure 4-2: Illustration for Case 1.1 for $n = 4$ and f is outside T but not a Star_n -center.




Figure 4-3: Illustration for Case 1.2 for $n = 4$ (before modification)




Figure 4-4: Illustration for Case 1.2 for $n = 4$ (after modification)

Section 4.2. Prime factorization of networks with respect to Δ_n

This section proves necessary and sufficient conditions for a Δ_n -partition being maximum and two structure theorems which characterize the Δ_n -poles, Δ_n -roots, Δ_n -zeroes and Δ_n -infinities of any given graph.

Theorem 4.2.1. *Let $(x_0, x_1, \dots, x_{2k})$ be a Δ_n -alternating path with respect to a maximum Δ_n -partition M . If x_0 is exposed, then each x_{2j} , $0 \leq j \leq k$, is a Δ_n -pole.*

Proof. By definition, x_0 is clearly a Δ_n -pole. We next prove that each x_{2k} , $1 \leq j \leq k$, is also a Δ_n -pole.

Consider a new Δ_n -partition $M' = M - (x_{2k-1}, x_{2k})$, that is, M' is formed from M by deleting the edge (x_{2k-1}, x_{2k}) . Hence M' is not a maximum Δ_n -partition and $(x_0, x_1, \dots, x_{2k-1})$ is a Δ_n -augmenting path with respect to the new Δ_n -partition M' . By Lemma 4.1.1, we can find another Δ_n -partition, say N , which covers x_0 and all other vertices covered by M except x_{2k} . In other words, x_{2k} is exposed by the maximum Δ_n -partition N . Thus x_{2k} is indeed a Δ_n -pole. ■

The next two theorems characterize the Δ_n -roots, Δ_n -poles and maximum Δ_n -partitions on any graph.

Theorem 4.2.2. *With respect to every non-maximum Δ_n -partition on a given graph, there exists at least one Δ_n -augmenting path.*

This theorem, together with Theorem 4.2.1, implies that a Δ_n -partition M on a given graph G is maximum if and only if M admits no Δ_n -augmenting path.

Theorem 4.2.3. (Structure theorem for Δ_n -partition) *For a given graph G , let P denote the set of all Δ_n -poles, R the set of all Δ_n -roots. Then,*

- (a) *The subgraph $G - (P \cup R)$ is a Δ_n -regular graph;*
- (b) *Every Δ_n -pole is adjacent to only Δ_n -roots. This, implies that every connected component of the subgraph of G induced on P is a single vertex, and together with the next statement, implies that every Δ_n -pole is a Δ_n -infinity;*
- (c) *R is a primary Δ_n -factorizer of G .*
- (d) *Every Δ_n -root is a Δ_n -zero which is adjacent to at least n Δ_n -poles. Moreover, there is a one-to- $(n-1)$ mapping from Δ_n -roots to their adjacent Δ_n -poles.*
- (e) *Let F be the induced subgraph of G on $P \cup R$. Then every vertex from P (resp. from R) is a Δ_n -pole (resp. Δ_n -root) of the graph F . Moreover, $\dim(G, \Delta_n) = \dim(F, \Delta_n)$.*

Proof of Theorem 4.2.2 and Theorem 4.2.3. The proof is trivial when G is Δ_n -regular, so we only

consider the case when G is Δ_n -singular. Consider any Δ_n -partition M on G such that there is no Δ_n -augmenting path in G with respect to M , and let z_1, z_2, \dots, z_d denote the vertices exposed by M . Apply Algorithm 4.1.2 on G with respect to M . Since M does not admit any Δ_n -augmenting path in G , Algorithm 4.1.2 can only terminate in Case 2. It can be easily checked that, at any time t , the following 5 basic properties are satisfied:

- (1) For every even vertex, there exists a unique Δ_n -alternating path in T_t with respect to M from an exposed vertex by M_t to that vertex.
- (2) A vertex x is outside T_t if and only if the whole class of x defined by M_t is outside T_t .
- (3) A vertex in T_t is a Star_n -center with respect to M if and only if it is an odd vertex.
- (4) The number of even vertices in T_t exceeds $n-1$ times the number of odd vertices by exactly d .
- (5) Every even vertex in T_t is adjacent to only odd vertices.

It follows from the fact that M exposes d vertices that $\dim(G, \Delta_n) \leq d$. On the other hand, let R'_τ denote the set of odd vertices in T_τ . From Property (4), there are at least $d+(n-1)|R'_\tau|$ even vertices. It then follows that $\dim(G, \Delta_n) \geq d$. By the same argument, we can prove that if x is a vertex in $G-V(T_t)$ (then x is necessarily a vertex in G), then $\dim(G-x, \Delta_n) \geq d$. We therefore reach the following conclusions:

- (6) $\dim(G_n, \Delta_n) = d$.
- (7) M is a maximum Δ_n -partition on G .
- (8) If x is a vertex in $G-V(T_t)$, then x can not be a Δ_n -pole of G .
- (9) The set R'_τ of odd vertices in T_τ is a Δ_n -factorizer of G .
- (10) Every even vertex is a Δ_n -pole of G , by Theorem 4.2.1.

Note that Property (7) offers a proof of Theorem 4.2.2. From Properties (8) and (10), we conclude that

(11) A vertex is a Δ_n -pole of G if and only if it is an even vertex in T_τ .

Moreover, by Property (3), we have

(12) A vertex is a Δ_n -root of G if and only if it is an odd vertex in T_τ , thus $R'_\tau = R$.

Statement (a) of Theorem 4.2.3 follows from Properties (11), (12) and (2). Statement (b) follows from Properties (5) and (10). Statements (d) and (d) follow from Properties (9), (12), (3) and (11) and Lemma 2.1.3.

Now we turn to prove Statement (e). From (a), the graph $G-V(F)$ is Δ_n -regular. Thus

$$d = \dim(G, \Delta_n) \leq \dim(G-V(F), \Delta_n) + \dim(F, \Delta_n) = \dim(F, \Delta_n).$$

On the other hand, M induces a Δ_n -partition on F which exposes d vertices, that is, $\dim(F, \Delta_n) \leq d$.

Therefore, we have

(13) $\dim(F, \Delta_n) = \dim(G, \Delta_n) = d$.

Let x be a vertex in P , that is, $\dim(G-x, \Delta_n) = d-1$. Consider a maximum Δ_n -partition M^* on G which exposes the vertex x . From Properties (11) and (12), there exist no Δ_n -poles of G in $G-V(F)$, thus the M^* induces a Δ_n -partition on F which exposes x and some other $d-1$ vertices. Thus

$$\dim(F-x, \Delta_n) \leq \dim(G-x, \Delta_n) = d-1 = \dim(F, \Delta_n)-1,$$

which implies that x is also a Δ_n -pole of F , so we have shown:

(14) $\dim(F, \Delta_n) = \dim(G, \Delta_n) = d$, and every vertex from P is a Δ_n -pole of the graph F .

From Properties (9) and (11) and Statement (a), R is a Δ_n -factorizer of F . Thus every vertex in R is a Δ_n -zero of F , by Lemma 2.1.3. Therefore every vertex in R is a Δ_n -root of F . Statement (e) then follows from Properties (13) and (14). ■

The following theorem is the Δ_n -counterpart to Mendelsohn-Dulmage Theorem. It follows directly from Lemma 4.1.1 and Theorem 4.2.2.

Theorem 4.2.4. *For a Δ_n -partition M on a given graph, there exists a maximum Δ_n -partition which*

covers all the vertices covered by M . In particular, for any given vertex, there exists a maximum Δ_n -partition covering that vertex.

The following theorem is a Δ_n -partition counterpart of Theorem 1.3.13, the Berge formula for the matching theory.

Theorem 4.2.5. *Let G be a graph. For any vertex subset S , denote by $i(G-S)$ the number of exposed vertices in the subgraph $G-S$. Then*

$$\dim(G, \Delta_n) = \max_S \{i(G-S) - (n-1)|S|\}.$$

Proof. For any vertex subset S of G , any vertex in S can only “save” at most $n-1$ exposed vertices.

We thus have for any vertex subset S ,

$$i(G-S) - (n-1)|S| \leq \dim(G, \Delta_n).$$

On the other hand, by Theorem 4.2.9, the set R of all Δ_n -roots is a primary Δ_n -factorizer of G .

Moreover, by The following theorem characterizes all Δ_n -singular Δ_n -prime graphs.

Theorem 4.2.12. *each Δ_n -singular connected component in $G-R$ is a single-vertex graph. Thus the equality holds when $S = R$, which establishes the theorem. ■*

It immediately follows from Theorem 4.2.5 that

Theorem 4.2.6. *A graph G admits a perfect Δ_n -partition if and only if for every vertex subset S of G , $i(G-S) \leq (n-1)|S|$, where $i(G-S)$ is the number of isolated vertices in the subgraph $G-S$.*

Theorem 4.2.3 characterizes Δ_n -roots and Δ_n -poles of any given graph. We next characterize Δ_n -zeros and Δ_n -infinities. We need the following lemma, which gives an equivalent definition of Δ_n -zero.

Lemma 4.2.7. *A vertex z in a graph G is a Δ_n -zero if and only if z is a $Star_n$ -center with respect to every maximum Δ_n -partition on G .*

Proof. (\Rightarrow) Assume, for contradictions, that there exists a maximum Δ_n -partition M of G such that

z is not a Star_n -center. Let X denote the class of z defined by M . Then $\{M \setminus X, \{z\}\}$, all connected components of the induced subgraph on $X \setminus \{z\}$ is a new Δ_n -partition which has singleton classes strictly less than $\dim(G, \Delta_n) + (n-1)$. Therefore,

$$\dim(G - z, \Delta_n) < \dim(G, \Delta_n) + (n-1),$$

that is, z is not a Δ_n -zero, which is a contradiction.

(\Leftarrow) Let M be an arbitrary maximum Δ_n -partition on G with the class of z . Let $X = \{z, x_1, x_2, \dots, x_{n-1}\}$ be the class of z . Since the induced graph on X is a Star_n with z as its center, we deduce that $M^* = (M \setminus X) \cup \{x_1\} \cup \{x_2\} \cup \dots \cup \{x_{n-1}\}$ is another Δ_n -partition on $G - z$. If M^* is not a maximum Δ_n -partition on $G - z$, then by Theorem 4.2.2, we can find a maximum Δ_n -partition M^{**} (on $G - z$) which covers a non-empty subset Y of $\{x_1, x_2, \dots, x_{n-1}\}$, as well as all the vertices covered by M^* . Then $M^{**} \cup (X \setminus Y)$ is a maximum Δ_n -partition on G under which the class $\{X \setminus Y\}$ of z is not a Star_n , a contradiction. Therefore M^* is a maximum Δ_n -partition on G , and hence

$$\dim(G - z, \Delta_n) = \dim(G, \Delta_n) + (n-1),$$

that is, z is a Δ_n -zero. ■

Theorem 4.2.8. *Given a graph G , let T denote the set of all Δ_n -infinities, Z the set of all Δ_n -zeros.*

Then,

(a) *The subgraph $G - (T \cup Z)$ is a Δ_n -regular graph.*

(b) *There is a 1-to-($n-1$) mapping from Δ_n -zeros to their adjacent Δ_n -infinities. Moreover,*

$$|T| = (n-1)|Z| + \dim(G, \Delta_n).$$

(c) *Let D be the induced subgraph of G on all edges that are incident to Δ_n -infinities. Then every vertex in T (resp. in Z) is a Δ_n -infinity (resp. Δ_n -zero) of the graph D . Moreover,*

$$\dim(G, \Delta_n) = \dim(D, \Delta_n).$$

Proof. Let M be a maximum Δ_n -partition on G . By Lemma 4.2.7, all Δ_n -zeros are centers of certain Star_n 's with respect to M . Let H be the subgraph of G which consists of all such Star_n 's in G . Randomly select an edge (e, f) outside H which is incident to a non- Star_n -center f in H . Let the edge be (e, f) , f be the vertex in H whose class is centered at a different vertex g .

We now prove that e is a Δ_n -zero. If, on the contrary, e is not a Δ_n -zero, then, by Lemma 4.2.7, we can find a maximum Δ_n -partition, say N , on G such that the class of e , denoted by E , is not a Star_n . Thus the induced subgraph on $E \cup \{f\}$ is Δ_n -regular. In other words, we can find a maximum Δ_n -partition N on G such that the class of g is not a Star_n , which is a contradiction since g is a Δ_n -zero. So e must be Δ_n -zero, that is, a Star_n -center in G .

It then follows that non- star_n -centers in H are Δ_n -infinities since they are only adjacent to Δ_n -zeros, and, by Lemma 4.2.7, they are not Δ_n -zeros. In addition, the exposed vertices by M are Δ_n -poles, by Theorem 4.2.3, they are Δ_n -infinities as well.

On the other hand, let x be any Δ_n -infinity which is not exposed by M . By definition, a Δ_n -infinity can only be adjacent to Δ_n -zeros, thus under M the class of x must be a Star_n and the Star_n -center must be a Δ_n -zero. Therefore, the set of Δ_n -infinities of G is the union of exposed vertices defined by M and the non- Star_n -centers in H . Consequently, the graph D in the theorem is the graph H plus the edges in G that are adjacent to non- Star_n -centers in H or the exposed vertices of M .

All the statements then immediately follow. ■

Theorem 4.2.9. *For any graph G , every primary Δ_n -factorizer contains all Δ_n -roots. Thus, the set of Δ_n -roots is the unique minimal primary Δ_n -factorizer.*

Proof. It is similar to that of Theorem 3.2.17, thus omitted. ■

Theorem 4.2.10. *A graph G is a Δ_n -prime graph if and only if there is no Δ_n -zero in G .*

Proof. The “if” part: it follows from Lemma 2.1.3.

The “only if” part: It is equivalent to prove that if there is at least one Δ_n -zero then G is not a Δ_n -prime graph. In fact, let z be a Δ_n -zero of G and $G_1 = G - z$. It follows from $\dim(G_1, \Delta_n) = \dim(G, \Delta_n) + (n-1) > 1$ that G_1 is a Δ_n -singular graph with dimension strictly greater than 1. Let R^* be the set of all Δ_n -roots of G_1 , then all components of $G_1 - R^*$ are either exposed vertices or Δ_n -regular graphs and

$$\dim(G_1 - R^*, \Delta_n) = \dim(G_1, \Delta_n) + (n-1)|R^*|.$$

In other words, all of the components in $G - (\{z\} \cup R^*)$ are either exposed vertices or Δ_n -regular subgraphs and

$$\begin{aligned} \dim(G - (\{z\} \cup R^*), \Delta_n) &= \dim(G - z, \Delta_n) + (n-1)|R^*| \\ &= \dim(G, \Delta_n) + (n-1) + (n-1)|R^*| \\ &= \dim(G, \Delta_n) + (n-1)|\{z\} \cup R^*| \end{aligned}$$

Thus $\{z\} \cup R^*$ is a primary Δ_n -factorizer of G and the graph G is not a Δ_n -prime graph. ■

Compared to the templates X_2, Δ_2 , prime factorizers of a graph G with respect to Δ_n ($n \geq 3$) can be easily characterized by the following theorem.

Theorem 4.2.11. *Let Z be the set of Δ_n -zeros in a graph G , then Z is the unique prime Δ_n -factorizer of G .*

Proof. By Theorem 4.2.8, Z is a primary Δ_n -factorizer of G . If one of the Δ_n -regular connected components of $G - Z$ is not a Δ_n -prime graph with S as one of its Δ_n -factorizers, then, by Lemma 2.1.5, $S \cup Z$ is a Δ_n -factorizer of G . From Lemma 2.1.3, we know that the vertices in S are Δ_n -zeros, which is a contradiction. Thus Z is a prime Δ_n -factorizer of G .

To prove the uniqueness of this prime Δ_n -factorizer, let F be a prime Δ_n -factorizer of G . By

Lemma 2.1.3, we have $F \subset Z$. If F is not equal to Z , one then checks that vertices in $Z \setminus F$ are Δ_n -zeros of the graph $G - F$. Thus, by Theorem 4.2.10, $G - F$ is not Δ_n -prime, which means that F is not a prime Δ_n -factorizer of G . ■

The following theorem characterizes all Δ_n -singular Δ_n -prime graphs.

Theorem 4.2.12. (Characterization of Δ_n -blossom) *A given connected graph is a Δ_n -singular Δ_n -prime graph if and only if it is the single-vertex graph.*

Proof. The “if” part: it is immediate.

The “only if” part: Assume that G is a connected Δ_n -singular graph of order strictly larger than one. Then there exists a vertex that is adjacent to a Δ_n -pole. From Theorem 4.2.3, this vertex can’t be a Δ_n -pole and hence must be a Δ_n -root, which is contradiction by Theorem 4.2.9. ■

Theorem 4.2.13. *For a vertex v in a given graph G , the following statements are true:*

- (a) *If under some maximum Δ_n -partition on G , v is not the center of its class, then the Δ_n -order of v is equal to zero, i.e., $\dim(G - v, \Delta_n) - \dim(G, \Delta_n) = 0$;*
- (b) *If the Δ_n -order of v is 1, then we can find a maximum Δ_n -partition on G such that the class of v is $Star_2$;*
- (c) *The Δ_n -order of v is larger than or equal to k , $2 \leq k \leq n-1$, if and only if for any maximum Δ_n -partition on G the class of v is a $Star_m$, $k+1 \leq m \leq n$, with v as its center.*

Proof. Statement (a) is clear. Statement (b) can be proved in the same way as the sufficiency part of Statement (c):

Sufficiency part for Statement (c): Assume, on the contrary, that the Δ_n -order t of v is strictly less than k . Let M be a maximum Δ_n -partition on G . Then, the class X of v is a $Star_m$, $k+1 \leq m \leq n$ with v as its center. Let $X = \{v, x_1, x_2, \dots, x_{m-1}\}$. Then, $M^* = \{M \setminus X, \{x_1\}, \{x_2\}, \dots, \{x_{m-1}\}\}$ is a Δ_n -partition on $G - v$. Since $m-1 \geq k > t$, M^* is not a maximum Δ_n -partition on $G - v$, and hence it can

be augmented to be a maximum Δ_n -partition, say M^{**} , such that at least t vertices in $X \setminus v$, say x_1, x_2, \dots, x_t , are exposed by M^{**} . Finally the classes in M^{**} and $\{v, x_1, x_2, \dots, x_t\}$ yield a maximum Δ_n -partition on G such that the class of v is a Star_{t+1} centered at v , which is contradiction.

Necessity part for Statement (c): Arbitrarily take a maximum Δ_n -partition M on G . Let X denote the class of v under M . To prove that the induced graph on X is a Star_m with v as its center, for some m ($k+1 \leq m \leq n$) we consider the following possible cases.

It follows from $\dim(G-v, \Delta_n) \geq \dim(G, \Delta_n) + k$ ($2 \leq k \leq n-1$) that $|X| \geq k+1 \geq 3$. If $(X = K_3)$ or $(X = \text{Star}_m$ ($k+1 \leq m \leq n$) and v is not the center vertex), then $\dim(G-v, \Delta_n) = \dim(G, \Delta_n)$. Hence v must be the center vertex of $X = \text{Star}_m$, $k+1 \leq m \leq n$. ■

It should be pointed out that when $k = n-1$, this theorem reduces to Lemma 4.2.7.

Theorem 4.2.14. *If for any vertex v of a given graph G , $\deg(v) \leq n-2$, then G is a Δ_n -prime graph.*

Proof. Since $\deg(v) \leq n-2$ for every vertex v in G , no vertex in G is Δ_n -zero, by Lemma 4.2.7. Thus, G must be a Δ_n -prime graph, by Theorem 4.2.10. ■

Theorem 4.2.15. *Let G be a connected graph. Then,*

(a) *If the degree of every vertex is strictly less than n , then G is a Δ_n -regular graph. Furthermore, if $|G| \geq n+1$, then G is a Δ_n -prime graph.*

(b) *If the degree of every vertex is larger than or equal to $|G|/n$, then G is a Δ_n -regular graph. Furthermore, if the degree of every vertex is strictly larger than $|G|/n$, then G is a Δ_n -prime graph. In particular, any complete graph of order k , $k \geq 2$, is a Δ_n -regular Δ_n -prime graph.*

Proof. (a) Since $\deg(v) \leq n-1$ for any vertex v in G , by Statement (d) of Theorem 4.2.3, there exist no Δ_n -roots in G . Furthermore, since G is a connected graph, there is no Δ_n -pole in G . It then follows from (a) of Theorem 4.2.3 that G is indeed a Δ_n -regular graph.

We next prove that if, furthermore, $|G| \geq n+1$, there does not exist Δ_n -zero in G . If, on the

contrary, a vertex z is a Δ_n -zero in G , then by (b) of Theorem 4.2.8, z must be adjacent to at least $n-1$ Δ_n -infinities, say $\{y_1, y_2, \dots, y_{n-1}\}$. Since $\deg(z) \leq n-1$, no other vertex is adjacent to z . Because $|G| \geq n+1$, there is at least another vertex, say x , which is adjacent to one of these Δ_n -infinities, say y_1 . By the definition of Δ_n -infinity, x must be a Δ_n -zero in G . Consider a Δ_n -partition under which the class of z is $X = \{z, y_1, y_2, \dots, y_{n-1}\}$. By Theorem 4.2.4 and Lemma 4.2.7, there is a maximum Δ_n -partition on G such that the class of z is still $X = \{z, y_1, y_2, \dots, y_{n-1}\}$ and the class of x is a Star_n with x as its center. But this implies that $\deg(x) \geq n$, which is a contradiction.

(b) We first prove that G is a Δ_n -regular graph if $\deg(v) \geq |G|/n$ for every vertex v . Let M be a maximum Δ_n -partition on G . Assume, for contradictions, that G is not Δ_n -regular. Let $\{x\}$ be a singleton class defined by M , and let y_1, y_2, \dots, y_k denote all the vertices adjacent to x , where $k = \deg(x) \geq |G|/n$. From Theorem 4.2.3, y_1, y_2, \dots, y_k are all Δ_n -roots. This implies that $|G| \geq 1+kn \geq 1+|G|$, which is absurd. So G must be a Δ_n -regular graph.

We next prove that if, furthermore, $\deg(v) \geq |G|/n+1$ for every vertex v , G is a Δ_n -prime graph. For an arbitrary vertex v of G , consider the graph $G' = G-v$. Note that the order of G' is $|G|-1$ and for every vertex u in G' ,

$$\deg(u) \geq [|G|/n+1]-1 \geq [|G|-1]/n = |G'|/n,$$

which implies that G' must be a Δ_n -regular graph. So, v can't be a Δ_n -zero of G . In other words, there is no Δ_n -zero in G , and thus G is a Δ_n -prime graph. ■

Theorem 4.2.16. For any positive integer m , let P_m, T_m, R_m and Z_m denote the sets of Δ_m -poles, Δ_m -infinities, Δ_m -roots and Δ_m -zeros of a given graph G with $|G| = N$, respectively. Then,

(a)

$$\emptyset = P_N \subseteq T_N \subseteq P_{N-1} \subseteq T_{N-1} \subseteq P_{N-2} \subseteq \dots \subseteq P_3 \subseteq T_3 \subseteq P_2$$

(b)

$$\emptyset = R_{N+1} = Z_{N+1} \subseteq R_N \subseteq Z_N \subseteq R_{N-1} \subseteq Z_{N-1} \subseteq \dots \subseteq R_3 \subseteq Z_3 \subseteq R_2$$

Proof. We will only prove Statement (a), since (b) can be similarly proven. It can be easily checked that G is Δ_N -regular, and hence $P_N = \emptyset$. Since $\deg(v) \leq N-1 = (N+1)-2$ for every vertex v , we deduce that $R_{N+1} = Z_{N+1} = \emptyset$, by Theorem 4.2.14.

The relationship $P_m \subseteq T_m$ immediately follows from tatement (b) of Theorem 4.2.3. Thus we need only to prove $T_m \subseteq P_{m-1}$ for $m \geq 3$.

Let x be a Δ_m -infinity, and let M be a maximum Δ_m -partition on G which covers x . If x is an exposed vertex of G , then $x \in P_{m-1}$ follows immediately. Otherwise, starting from x , we shall iteratively construct a graph E whose vertices are labeled as even vertices or odd vertices.

Initially, E contains only vertex x , which is labeled as an *even vertex*. We next grow E by adding vertices and edges in G .

By the definition of Δ_m -infinity, the vertices adjacent to x are all Δ_m -zeros. For one of such Δ_m -zeros, say y , its class under M is a Star_m centered at y , by Lemma 4.2.7. We now extend the graph E by adding the edge (x, y) and the star shaped graph induced on the whole class of y . Among the new vertices of E , the vertex y is labeled as an *odd vertex*, while the other vertices are labeled as *even vertices*.

In general, let z be an even vertex in E . If no vertices outside E are adjacent to z , then we turn to consider the other even vertices in E . If g is a vertex adjacent to z but not a vertex of E , then the class of g defined by M is a Star_m centered at g . (If this is not the case we can find a maximum Δ_m -partition on G such that the class of x is no more at Star_m , which is contradictory to the assumption that x is a Δ_m -zero.) We then extend the graph E by adding the edge (g, z) and the star shaped graph induced on the whole class of g . Among the new vertices of E , the vertex g is labeled an odd vertex and the other vertices are labeled even vertices.

Upon termination of this iterative process, one observes that in E , all the classes defined by M

are Star_m 's. Deleting exactly one non-center vertex from each of these Star_m 's, we then obtain a group of Star_{m-1} 's. Let M' be a Δ_{m-1} -partition on G with the above Star_{m-1} 's as its classes. Then through the augmenting procedure in the proof of Theorem 4.2.3, we obtain a maximum Δ_{m-1} -partition on G which does not cover vertex x (Star_{m-1} classes are not changed in the procedure; in particular, the class of vertices adjacent to x are not changed). Thus x is a Δ_{m-1} -pole, and (a) then follows. ■

Chapter 5. Prime Factorization with Respect to the Template X_n

This chapter concentrates on network factorization theory with respect to X_n -partitions. Throughout this chapter, n stands for an arbitrary but fixed integer greater than or equal to 3.

We first restate some basic definitions and notation about X_n -partition.

- ◆ An X_n -partition of a graph G divides the vertices of G into classes such that the induced subgraph on each class is isomorphic to Star_k , $1 \leq k \leq n$.
- ◆ Given a graph G , the minimum number of vertices exposed by an X_n -partition is called the X_n -dimension of G , denoted by $\dim(G, X_n)$.
- ◆ When $\dim(G, X_n) > 0$, we say that the graph G is X_n -singular. When $\dim(G, X_n) = 0$, we say that G is X_n -regular.
- ◆ If a particular X_n -partition defines exactly $\dim(G, X_n)$ singletons, that partition is called a maximum X_n -partition on G . In the case of an X_n -regular graph, a maximum X_n -partition does not expose any vertex and is therefore called a perfect X_n -partition.
- ◆ The X_n -order of a vertex x in a graph G is defined as $\dim(G-x, X_n) - \dim(G, X_n)$. A vertex with X_n -order -1 is called an X_n -pole. A necessary and sufficient condition for a vertex to be an X_n -pole is it being exposed by a maximum X_n -partition.
- ◆ The maximum X_n -order of any vertex is $\Phi(X_n) = n-1$. A vertex with X_n -order $n-1$ is called an X_n -zero.
- ◆ If a vertex is not an X_n -pole vertex but is adjacent to at least one X_n -pole, then it is called an X_n -root.
- ◆ A vertex is called an X_n -infinity if all adjacent vertices (if any) are X_n -zeroes.
- ◆ A set S of vertices is called an X_n -factorizer of G if there are $\dim(G, X_n) + \Phi(X_n)|S|$ connected

components of $G-S$ that are X_n -singular graphs. Since $\Phi(X_n) = n-1$, this implies that exactly $\dim(G, X_n) + (n-1)|S|$ connected components of $G-S$ have the X_n -dimension 1 and all others are X_n -regular graphs.

- ◆ An X_n -prime graph is defined as a connected graph that has no non-empty X_n -factorizer.
- ◆ An X_n -factorizer S is called a primary X_n -factorizer if all connected X_n -singular components of $G-S$ are X_n -prime graphs (while the X_n -regular components may or may not be X_n -prime graphs.).
- ◆ An X_n -factorizer S of a graph is called a prime X_n -factorizer if all connected components of $G-S$ are X_n -prime graphs.

Section 5.1. An Edmonds-type algorithm

Let M be an X_n -partition on a graph G . A path on a graph is said to be an *alternating path* with respect to M if pairs of adjacent vertices on the path are alternately classmates and non-classmates under M . An X_n -alternating path (x_0, x_1, \dots, x_k) with respect to M is called an X_n -augmenting path with respect to M if the following conditions are satisfied:

1. x_0 is exposed by M .
2. If k is an odd integer, then the class of x_k is not a two-vertex class and x_k is not a Star_n -center with respect to M .
3. If k is an even integer, then for some m , $0 \leq m \leq k/2$, the vertex x_{2m+1} belong to a two-vertex class and x_k is adjacent to x_{2m+1} but not to x_{2m} .

Lemma 5.1.1. *Let (x_0, x_1, \dots, x_k) be an X_n -augmenting path with respect to an X_n -partition M . Then there exists an X_n -partition that covers x_0 and all vertices covered by M .*

Proof. The proof is by induction on k .

Case 1. k is odd.

If the class of x_{k-2} is not a two-vertex class and x_{k-2} is not a Star_n -center with respect to M , then $(x_0, x_1, \dots, x_{k-2})$ is a shorter X_n -augmenting path with respect to M . The induction hypothesis applies. Otherwise, we modify M to form a new X_n -partition by (reassigning the vertex x_{k-1} into the class of x_k , if the class of x_k is centered at x_k), or (replacing $x_{k-1} \cup \{\text{the class of } x_k \text{ under } M\}$ by two new classes $\{x_{k-1}, x_k\}$ and $\{\text{the class of } x_k \text{ under } M\} - \{x_k\}$, if x_k is not the center vertex of its class under M). In either case, $(x_0, x_1, \dots, x_{k-2})$ is a shorter X_n -augmenting path with respect to the new X_n -partition, which covers all vertices covered by M . The induction hypothesis then applies.

Case 2. k is even.

If for some integer i , $i < k$, such that the segment (x_0, x_1, \dots, x_i) is a shorter X_n -augmenting path with respect to X_n -partition M , the induction hypothesis applies. We assume the following two conditions in the remaining proof:

(1) there exists some integer m , $0 \leq m < k/2$, such that the vertex x_{2m+1} belongs to a two-vertex class under M and x_k is adjacent to x_{2m+1} but not to x_{2m} ; and

(2) for every i , the vertex x_{2i+1} either belongs to a two-vertex class or is a Star_n -center.

We will modify the X_n -partition M to obtain a new X_n -partition N , which covers all vertices covered by M . Let $C_M(x_j)$ denote the class of x_j under M . The modification is to replace the classes of all $x_{2m+1}, x_{2m+2}, \dots, x_k$ by the following new classes:

$$\{x_{2m+1}, x_k\}, \{C_M(x_{2m+3}) \cup \{x_{2m+2}\} - \{x_{2m+4}\}\}, \\ \{C_M(x_{2m+5}) \cup \{x_{2m+4}\} - \{x_{2m+6}\}\}, \dots, \{C_M(x_{2j+1}) \cup \{x_{2j}\} - \{x_{2j+4}\}\}, \dots, \{C_M(x_{k-1}) \cup \{x_{k-2}\} - \{x_k\}\}$$

It can be checked that the path $(y_0 = x_0, y_1 = x_1, \dots, y_{2m+1} = x_{2m+1}, y_{2m+2} = x_k)$ is a shorter X_n -augmenting path with respect to the new X_n -partition N . The induction hypothesis applies. Figure 5-1 and Figure 5-2 show an example of this modification from M to N , where we assume that $n = 4$, $k = 8$ and $m = 1$. ■




Figure 5-1: Before modification from M to N for $n = 4$, $k = 8$ and $m = 1$.




Figure 5-2: After modification from M to N for $n = 4$, $k = 8$ and $m = 1$, we have a new shorter augmenting path with respect to N.

A vertex x of a graph G is called a *cut-vertex* if the number of connected components of the graph $G-x$ is greater than that of the original graph G . A subgraph H of G is said to be a *block* of G if H is a maximal subgraph without any cut-vertex in H . A graph is called *KB-blossom* if each of its blocks is a complete graph of odd order. Clearly every KB-blossom is a blossom with X_n -dimension 1.




Figure 5-3 An example of KB-Blossom with two K_5 blocks and four K_3 blocks “glued” together

Lemma 5.1.2. *Let a and b be two adjacent vertices in a graph H . Assume that $G = H - \{a, b\}$ is a KB-blossom. Then exactly one of the following is true:*

- (a) *There is no edge between G and $\{a, b\}$;*
- (b) *There is a vertex of G that is adjacent to both a and b , and every other vertex of G is adjacent to neither a nor b . In this case, H is a KB-blossom;*
- (c) *There is a block in G such that all of its vertices are adjacent to both a and b and all vertices of G outside this block adjacent to neither a nor b . In this case, H is also a KB-blossom;*
- (d) *Given any maximum X_2 -partition M on H with a and b sharing a class, there exists an X_n -augmenting path with respect to the X_n -partition M . In this case, H is an X_n -regular graph.*

Proof. Let M be a maximum X_2 -partiton on H with a and b sharing a class, and e be the only exposed vertex of G by M .

If a vertex x in G is adjacent to a but not to b , then there exists an even-length alternating path in G from x to e . Gluing this path and the path (x, a, b) together, we have an even-length X_n -augmenting path with respect to the X_n -partition M . Similarly, if a vertex in G is adjacent to b but not to a , then there also exists an X_n -augmenting path with respect to the X_n -partition M .

Therefore we assume, hereafter, that if a vertex in G is adjacent to a , then it is also adjacent to b , and vice versa. Consider the following cases.

Case 1. At most one vertex in G is adjacent to a and b . Then Statements (a) or (b) stands.

Case 2. At least two vertices in G are adjacent to a and b .

Case 2.1. There exist two non-adjacent vertices, say x and y , in G such that they are adjacent to a and b . Let x' and y' denote the classmates of x and y under M , respectively. By Lemma 1.3.9, there

exists an even-length alternating path P (resp. Q) in G from e to x (resp. y). Obviously, either P is inside $G - \{y, y'\}$ or Q is inside $G - \{x, x'\}$; in this proof, we assume that P is inside $G - \{y, y'\}$. If y' is adjacent to neither a nor b , we glue the path P and the path $\{x, a, b, y, y'\}$ together to form an even-length X_n -augmenting path with respect to the X_n -partition M (Note that y' is adjacent to y but not b ; see Figure 5-4). If y' is adjacent to a and b , we glue the path P and the path $\{x, a, b, y', y\}$ together to form an even-length X_n -augmenting path with respect to the X_n -partition M (Note that y is adjacent to a but not to x ; see Figure 5-5).

Case 2.2. Those vertices in G that are adjacent to a and b all reside in a block B of G .

Case 2.2.1. Every vertex in B is adjacent to a and b . Then Statement (c) stands.

Case 2.2.2. Two classmates x and x' in B are both adjacent to a and b but the third vertex y in B does not. By Lemma 1.3.9, there exists an even-length alternating path P from e to y ; again, we can assume that P is in $G - \{x, x'\}$. Gluing the path P and the path $\{y, x, x', a, b\}$ together, we have an even-length X_n -augmenting path with respect to the X_n -partition M (Note that b is adjacent to x but not to y ; see Figure 5-6).

Case 2.2.3. x is adjacent to a and b but its classmate x' is not. Let y be another vertex that is adjacent to a and b . By Lemma 1.3.9, there exists an even-length alternating path P , in $G - \{x, x'\}$, from e to y . Gluing the path P and the path $\{y, a, b, x, x'\}$ together, we have an even-length X_n -augmenting path with respect to the X_n -partition M (Note x' is adjacent to x but not to b ; see Figure 5-7).




Figure 5-4: Illustration of Case 2.1.




Figure 5-5: Illustration of Case 2.1.




Figure 5-6: Illustration of Case 2.2.2.




Figure 5-7: illustrates Case 2.2.3.

Algorithm 5.1.3. Given an X_n -partition M on a graph G , this algorithm determines whether G admits an augmenting path with respect to M . Write $G_0 = G$ and $M_0 = M$. The algorithm will construct a sequence of graphs G_t , $0 \leq t \leq \tau$, with an X_n -partition M_t on each G_t . In the end, whether there is an augmenting path with respect to M_n in G_n will be apparent. If there is, then, for every t , an augmenting path with respect to M_{t+1} in G_{t+1} induces an augmenting path with respect to M_t in G_t . The graph G_t will be associated with, besides the matching M_t , an acyclic subgraph T_t , in which every vertex is labeled either *even* or *odd* so that T_t is a bipartite graph between even and odd vertices. Figure 3-2 illustrates G_t , M_t and T_t for a generic t .

Initially, those vertices z_1, z_2, \dots, z_d exposed by M are all labeled as even vertices. Let T_0 consist of these d vertices. Given G_t , M_t and T_t , the corresponding iterative step in the algorithm

achieves exactly one of the following:

- (a) Keep both G_t and M_t the same, whereas grow T_t by adding an odd vertex, $(n-1)$ even vertices, and n edges. The first edge is between an existing even vertex and the new odd vertex; the remaining $(n-1)$ edges, which belong to M_t , are between the new odd vertex and the new even vertices. At the end of this step, increase the index t by 1.
- (b) Contract a triangle in T_t (and G_t) to obtain T_{t+1} (and G_{t+1}), and let M_t induce an X_n -partition M_{t+1} on G_{t+1} . At the end of this step, increase the index t by 1.
- (c) Identify an augmenting path of M_t , and recursively find an augmenting path with respect to M . The algorithm terminates, that is, t is the final index τ .
- (d) G_τ does not admit any augmenting path with respect to M_τ , and hence G does not admit any augmenting path with respect to M . The algorithm terminates.

The iterative step at time t starts by looking for an edge of G_t such that it is

- not an edge of T_t ,
- incident to at least one even vertex of T_t , and
- is not incident to any odd vertex of T_t .

Case 1. Such an edge exists, say, the selected edge is (e, f) , where e is an even vertex of T_t and f is either an even vertex of T_t or outside T_t . We consider the following subcases:

Case 1.1. The vertex f is outside T_t .

Case 1.1.1. The class of f defined by M is a two-vertex class $\{f, y\}$ and y is adjacent to e in G_t . Necessarily, y is also outside T_t . We consider two subcases depending on whether the triangle $\{f, e, y\}$ in G_t corresponds to an X_n -regular graph or X_n -singular graph in G (It will be clear that each vertex in G_t corresponds to a KB-blossom in G later in the algorithm.).

Case 1.1.1.1. The triangle $\{f, e, y\}$ in G_t corresponds to an X_n -regular graph in G . Let B_e be the

KB-blossom corresponding to the even vertex e . M induces a maximum X_2 -partition, say M_e , on B_e with the only exposed vertex, say s_e . By Lemma 5.1.2, there exists an X_n -augmenting path, say P_1 , starting from the vertex s_e in B_e with respect to M_e . On the other hand, there is another alternating path, say P_2 , in G with respect to M from an exposed vertex of M to s_e . Gluing the paths P_1 and P_2 together, an X_n -augmenting path in G with respect to M is formed. In this case, the iterative process stops. See Figure 5-8. ((c) is achieved)

Case 1.1.1.2. The triangle $\{f, e, y\}$ in G_t corresponds to an X_n -singular graph in G . By Lemma 5.1.2, this subgraph must be a KB-blossom. The iterative step, in this case, is to transform G_t and T_t by contracting the triangle $\{f, e, y\}$ into a single vertex in T_{t+1} which is labeled as an even vertex. Meanwhile the M_t induces an X_n -partition M_{t+1} on T_{t+1} . See Figure 5-9, Figure 5-10 and Figure 5-11. ((b) is achieved)

Case 1.1.2. The class of f , defined by the X_n -partition M , is a $Star_n$ with f as its $Star_n$ -center. Then necessarily, all members in this $Star_n$ are outside T_t . In this case, we add the n vertices of the $Star_n$ into T_t to obtain T_{t+1} . Moreover, we add into X_n -partition M_t a new class that consists of the n new vertices of T_t to obtain M_{t+1} , and we label f as odd and the other $(n-1)$ vertices as even. The graph G_t is unchanged, namely, set $G_{t+1} = G_t$. See Figure 5-12 and Figure 5-13. ((a) is achieved)

Case 1.1.3. Otherwise, let t_e be the vertex in the KB-blossom in G corresponding to the even vertex e such that t_e is incident to f . Then, there is an alternating path in G with respect to M from an exposed vertex by M to the vertex t_e . Glue this path and the path (t_e, f) to form an X_n -augmenting path in G with respect to M . In this case, the iterative process stops. See Figure 5-14. ((c) is achieved)

Case 1.2. The vertex f is an even vertex in T_t . Let t_e and t_f be two adjacent vertices of G in the KB-blossoms corresponding to e and f , respectively. Then, M induces a maximum X_2 -partition, say

M_e , on the KB-blossom, say B_e , corresponding to e . The only exposed vertex by M_e is denoted by s_e . The class of s_e under M is clearly a Star_n and s_e is not the center vertex. Then, there exists an even length alternating path, say P_1 , in G with respect to M from an exposed vertex by M to the vertex t_f . By Lemma 1.3.9, there exists an even-length alternating path, say P_2 , in B_e with respect to M_e from t_e to s_e . Gluing the paths P_1 , (t_f, t_e) and P_2 , and an odd-length X_n -augmenting path in G with respect to M is formed. If this is the case, the iterative process stops. See Figure 5-15. ((c) is achieved)

Case 2. The opposite to Case 1 occurs, i.e., every edge of G_t that is incident to at least one even vertex and no odd vertex must be an edge of T_t . ((d) is achieved) ■




Figure 5-8 illustrates the Case 1.1.1.1 for $n = 3$.




Figure 5-9: Illustration of Case 1.1.1.2, for $n = 3$: before contraction.

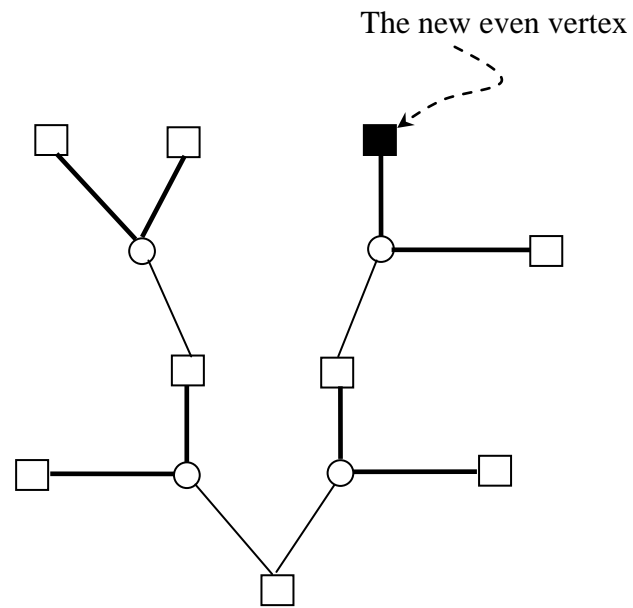


Figure 5-10: Illustration of Case 1.1.1.2, for $n = 3$: after contraction.




Figure 5-11: Illustration of Case 1.1.1.2, for $n = 3$: The KB-blossom corresponding to the new even vertex.




Figure 5-12: Illustration of Case 1.1.2, for $n = 3$: before modification.




Figure 5-13: Illustration of Case 1.1.2, for $n = 3$: after modification.




Figure 5-14: illustrates the Case 1.1.3 for $n=3$.




Figure 5-15: illustrates the Case 1.2 for $n = 3$.

Section 5.2. Prime factorization of networks with respect to X_n

The following theorem is the converse of Lemma 5.1.1. These two theorems together assert that an X_n -partition M is maximum on a graph if and only if M admits no X_n -augmenting path.

Theorem 5.2.1. *With respect to every non-maximum X_n -partition on a graph there exists an X_n -augmenting path.*

Theorem 5.2.1 will be proved together with the following structure theorem.

Theorem 5.2.2. *Let P and R denote the sets of X_n -poles and X_n -roots, respectively, of a graph G .*

Then,

- (a) *The subgraph $G - (P \cup R)$ is an X_n -regular graph;*

- (b) Every connected component of the induced subgraph on P is a KB-blossom;
- (c) R is a primary X_n -factorizer.
- (d) Every X_n -root is adjacent to at least $n-1$ such KB-blossoms;
- (e) Let F be the induced subgraph of G on $P \cup R$. Then every vertex from P (resp. R) is an X_n -pole (resp. X_n -root) of the graph E . Moreover $\dim(G, X_n) = \dim(F, X_n)$

Proof of Theorem 5.2.1 and Theorem 5.2.2: If G is an X_n -regular graph, the $P = R = \emptyset$. All statements are trivial. So, we assume that G is an X_n -singular graph. Consider any X_n -partition M on G such that there is no augmenting path in G with respect to M , and let z_1, z_2, \dots, z_d denote the vertices exposed by M . Apply Algorithm 5.1.3 on G with respect to M . Since M does not admit any X_n -augmenting path in G , Algorithm 5.1.3 can only terminate in Case 2. It can be easily checked that, at any time t , the following 5 basic properties are satisfied:

- (1) Every odd vertex in T_t is a vertex of the original graph G , so is every vertex in $G_t - V(T_t)$.
Every even vertex in T_t corresponds to a KB-blossom in G . Moreover, M induces a maximum X_2 -partition on every such KB-blossom.
- (2) If two vertices f and g belong to the same class defined by M , then either both f and g or neither of them are vertices in $G_t - V(T_t)$.
- (3) In T_t , a vertex is a Star_n -center defined by M_t if and only if it is an odd vertex.
- (4) Every connected component of T_t contains exactly one exposed vertex of M_t .
- (5) The number of even vertices in T_t exceeds $n-1$ times the number of odd vertices by exactly d .

We next deduce from the above five properties the sixth property:

- (6) In the original graph G , with respect to M , there exists an odd-length (resp. even-length) alternating path from an exposed vertex by M to an odd vertex in T_t (resp. a vertex in a

KB-blossom corresponding to an even vertex of T_t).

We shall only prove the odd-length part of (6), the other part being similar. Let x be an odd vertex in T_t . From Property (4), there exists a unique odd-length alternating path in T_t with respect to M_t that connects an exposed vertex by M_t to x . Let this path be $(x_0, x_1, x_2, \dots, x_{2n-1}, x_{2n}, x_{2n+1} = x)$. From Property (3), the vertices x_{2i-1} and x_{2i} belong to the same class of M_t for $1 \leq i \leq n$. For $0 \leq i \leq n$, let B_{2i} be the KB-blossom (also a blossom) in G corresponding to the even vertex x_{2i} , and s_{2i} the exposed vertex by M on B_{2i} , and t_{2i} a vertex in B_{2i} that is adjacent to the odd vertex x_{2i+1} . From Lemma 1.3.9, there exists an even-length alternating path in G with respect to M from s_{2i} to t_{2i} , for $0 \leq i \leq n$. These alternating paths and the paths $(t_{2i}, x_{2i+1}, s_{2i+2})$, $0 \leq i \leq n$, and (t_{2n}, x_{2n+1}) can be pieced together to form an odd-length alternating path in G with respect to M from the exposed vertex s_0 of M to x_{2n+1} .

Since T_t is bipartite, we conclude that

(7) In G_t , every even vertex is adjacent to only odd vertices.

Let R'_τ denote the set of odd vertices in T_τ . From Property (5), there are at least $d + |R'_\tau| (n-1)$ even vertices. From Property (7), each of these even vertices is by itself a connected component in $G - R'_\tau$. Thus, $\dim(G_\tau, X_n) \geq d$. On the other hand, let M^* be the X_n -partition on G_τ which coincide with the X_n -partition M_τ on T_τ and with the X_n -partition M on $G_\tau - V(T_\tau)$. Then M^* exposes exactly d vertices in T_τ and none in $G_\tau - V(T_\tau)$. Thus $\dim(G_\tau, X_n) \leq d$. We therefore reach the following conclusions:

(8) $\dim(G_\tau, X_n) = d$.

(9) M^* is a maximum X_n -partition on G_τ .

(10) Every even vertex is an X_n -pole of G_τ .

In fact, let e be an even vertex. From Properties (4) and (3), there exists in T_τ an alternating

path with respect to M_τ that connects an exposed vertex by M_τ to e . Since the X_n -partition M^* coincides with M_τ on T_τ , the alternating path with respect to M_τ is also an alternating path with respect to M^* and the exposed vertex of M_τ is also exposed by M^* . Thus there exists in G_τ an alternating path with respect to the maximum X_n -partition M^* that connects an exposed vertex of M^* to e . This proves that e is an X_n -pole of G_τ .

From Property (1), vertices in R'_τ , i.e., the odd vertices in T_τ , are also vertices of the original graph G and there are at least $d+|R'_\tau|(n-1)$ KB-blossom components in $G-V(R'_\tau)$. Thus $\dim(G, X_n) \geq d$. By the same argument, if x is a vertex in $G_\tau-V(T_\tau)$ (and hence also a vertex of G , by Property (1)), then $\dim(G-x, X_n) \geq d$. On the other hand, the X_n -partition M on G exposes exactly d vertices.

We therefore have the following conclusions:

$$(11) \dim(G, X_n) = d.$$

(12) If x is a vertex in $G_\tau-V(T_\tau)$, then x is not an X_n -pole of G .

(13) R'_τ is an X_n -factorizer of G .

(14) M is a maximum X_n -partition on G .

Thus Theorem 5.2.1 is proved. We next claim that

(15) A vertex of G is an X_n -pole if and only if it belongs to an KB-blossom in G that is contracted into an even vertex of T_τ .

From Property (13), every odd vertex of T_τ is an X_n -zero of G . This together with Property (12) proves the “only if” part of the above claim. Conversely, let B be a KB-blossom in G that is contracted into an even vertex e' of T_τ and let e be a vertex in B . We need to show that e is an X_n -pole of G . From Properties (8) and (10), we know that $\dim(G_\tau-e', X_n) = d-1$. Note that the X_n -dimension of a graph is unchanged when a KB-blossom in it is contracted into a single vertex

such that this vertex is an X_n -pole in the resulting graph. Therefore $\dim(G-B, X_n) = \dim(G_{\tau-e'}, X_n)$.

Thus

$$\dim(G-e, X_n) \leq \dim(G-B, X_n) + \dim(B-e, X_n) = d-1+0 = \dim(G, X_n)-1,$$

by Property (11). Property (15) is then proved.

By Properties (13) and (1), every odd vertex is an X_n -zero of G . Thus every vertex in R'_{τ} is an X_n -root of G , by Properties (3) and (15). On the other hand, by Properties (7) and (15), there exist no X_n -roots other than the odd vertices. Therefore we have proved that

(16) A vertex of G is an X_n -root if and only if it is an odd vertex of T_{τ} , i.e., $R'_{\tau} = R$.

(17) The induced subgraph of G on the vertices of $G_{\tau} - V(T_{\tau})$ is an X_n -regular graph.

Statement (a) is implied by Properties (15), (16) and (17). Statement (b) follows from by Properties (15) and (3). From Properties (13) and (16), R is an X_n -factorizer of G . Moreover, by Properties (7), (1) and (17), every X_n -singular connected component of $G-R$ is a KB-blossom, which is an X_n -prime graph. Statements (c) and (d) are then approved.

Now we begin to prove (e). From (a), the graph $G-F$ is X_n -regular. Thus,

$$d = \dim(G, X_n) \leq \dim(G-F, X_n) + \dim(F, X_n).$$

On the other hand, the X_n -partition M induces an X_n -partition on F which exposes d vertices, i.e., $\dim(F, X_n) \leq d$. Hence, we have

$$(18) \dim(F, X_n) = \dim(G, X_n) = d.$$

Let x be a vertex in P , i.e., $\dim(G-x, X_n) = d-1$. Consider a maximum X_n -partition M_1 on G which exposes the vertex x . From Properties (15) and (16), there exists no X_n -pole of G in $G-F$, thus the M_1 induces an X_n -partition on e which exposes x and some other $d-1$ vertices. Thus,

$$\dim(F-x, X_n) \leq \dim(G-x, X_n) = d-1 = \dim(F, X_n)-1,$$

which implies that x is also an X_n -pole of F , i.e.,

(19) Every vertex from P is an X_n -pole of the graph F .

From (a) and (c), R is also an X_n -factorizer of F . Thus every vertex in C is an X_n -zero of F . Therefore every vertex in R is an X_n -root of F . Statement (e) then follows from Properties (17) and (18). ■

The following theorem is an X_n -partition counterpart of Theorem 1.3.13, the Berge formula for the matching theory.

Theorem 5.2.3. *Let G be a graph. For any vertex subset S of G , denote by $p(G-S)$ the number of connected components in the subgraph $G-S$ that are KB-blossoms. Then*

$$\dim(G, X_n) = \max_S \{p(G-S) - (n-1)|S|\}.$$

Proof. For any vertex subset S of G , any vertex in S can only “save” at most $n-1$ vertices. We thus have for any vertex subset S ,

$$p(G-S) - (n-1)|S| \leq \dim(G, X_n).$$

By Theorem 5.2.8, the set R of all X_n -roots is a primary X_n -factorizer of G . Moreover, by Theorem 5.2.2, each X_n -singular connected component in $G-S$ is a KB-blossom. Thus, Thus the equality holds when $S = R$, the set of Δ_n -roots of G . ■

It immediately follows from Theorem 5.2.3 that

Theorem 5.2.4. *A graph G admits a perfect X_n -partition if and only if for every vertex subset S of G , $p(G-S) \leq (n-1)|S|$, where $p(G-S)$ is the number of connected components in the subgraph $G-S$ that are KB-blossoms in the subgraph $G-S$.*

The following theorem is the X_n -partition counterpart to Mendelsohn-Dulmage Theorem. It follows directly from Lemma 5.1.1 and Theorem 5.2.1.

Theorem 5.2.5. *Let P be an X_n -partition on G . Then there exists a maximum X_n -partition M which*

covers all the vertices of G covered by P . In particular, for a given vertex in G , there is always a maximum X_n -partition which covers that vertex.

Proof. This corollary follows immediately from Lemma 5.1.1 and Theorem 5.2.1. ■

The following theorem characterizes all X_n -singular X_n -prime graphs.

Theorem 5.2.6. (Characterization of X_n -blossom) *The following statements are equivalent for a graph G .*

- (a) *G is a KB-blossom.*
- (b) *G is an X_n -prime graph with X_n -dimension equal to 1.*
- (c) *G is an X_n -singular X_n -prime graph, i.e. an X_n -blossom.*
- (d) *G is connected and all vertices are X_n -poles.*

Proof. (1) \Rightarrow (2). Due to Lemma 3.2.10 and the fact that every KB-blossom has an X_n -dimension equal to 1.

(2) \Rightarrow (3). Obvious.

(3) \Rightarrow (4). G is X_n -singular implies the existence of X_n -pole. If there is at least one vertex that is not X_n -pole, by the connectedness of G , there exists at least one X_n -root which is contradictory to the assumption that G is X_n -prime.

(4) \Rightarrow (1). Due to Theorem 5.2.2. ■

Lemma 5.2.7. *A vertex z is an X_n -zero of a graph G if and only if z is a $Star_n$ -center with respect to every maximum X_n -partition on G .*

Proof. (\Rightarrow) Let z be an X_n -zero of G . If, on the contrary, there exists a maximum X_n -partition M on G such that the class X of z defined by M is not $Star_n$ with z as its center. Then the partition $\{M \setminus X, \text{the connected components on the induced subgraph on } X - z, \{z\}\}$ is a new X_n -partition on

G which has fewer than $\dim(G, X_n) + (n-1)$ singleton classes. Therefore,

$$\dim(G-z, X_n) < \dim(G, X_n) + (n-1),$$

which implies that z is not an X_n -zero, a contradiction.

(\Leftarrow) Let M be an arbitrary maximum X_n -partition on G under which the class of z is X . Then X is a Star_n , say $X = \{z, x_1, x_2, \dots, x_{n-1}\}$. One checks that

$$M^* = (M \setminus X) \cup \{x_1\} \cup \{x_2\} \cup \dots \cup \{x_{n-1}\}$$

is another X_n -partition on $G-z$. If M^* is not a maximum X_n -partition on $G-z$, then by Theorem 5.2.1 we can find a maximum X_n -partition M^{**} such that M^{**} covers at least one of $\{x_i\}$, $1 \leq i \leq n-1$, as well as all vertices covered by M^* . Let Y be the subset of all vertices in $\{x_1, x_2, \dots, x_{n-1}\}$ covered by M^{**} . Then $M^{**} \cup (X \setminus Y)$ is a maximum X_n -partition on G under which the class $\{X \setminus Y\}$ of z is not a Star_n , which is contradictory to the sufficiency assumption. Therefore M^* is a maximum X_n -partition on G , hence

$$\dim(G-z, X_n) = \dim(G, X_n) + (n-1),$$

i.e., z is an X_n -zero. ■

Theorem 5.2.8. *For any graph G , every primary Δ_n -factorizer contains all Δ_n -roots. Thus, the set of Δ_n -roots is the unique minimal primary Δ_n -factorizer.*

Proof. It is similar to that of Theorem 3.2.17, thus omitted. ■

Theorem 5.2.9. *A graph G is an X_n -prime graph if and only if there is no X_n -zero in G .*

Proof. Similar to that of Theorem 4.2.10, thus omitted. ■

Theorem 5.2.10. *Let Z be the set of X_n -zeros in a graph G , then Z is the unique prime X_n -factorizer of G .*

Proof. It can be checked that Z is a primary X_n -factorizer of G . The remainder of the proof is then the same as that of Theorem 4.2.11. ■

Theorem 5.2.11. *For a vertex v in a graph G , we have*

- (a) *If under some maximum X_n -partition on G , the vertex v is not the center vertex of its class, then the X_n -order of v is equal to zero, i.e., $\dim(G-v, X_n) - \dim(G, X_n) = 0$;*
- (b) *If v is a vertex with its X_n -order equal to 1, then we can find a maximum X_n -partition on G such that the class of v is $Star_2$;*
- (c) *v is a vertex with its X_n -order larger than or equal to k , $2 \leq k \leq n-1$, if and only if for any maximum X_n -partition on G the class of v is a $Star_m$ with v as its center, $k+1 \leq m \leq n$.*

Proof. Statement (a) is clear. Statement (b) can be proved in the same way as the following sufficiency part of (c).

Sufficiency part of Statement (c): Assume, on the contrary, that the X_n -order t of v is strictly less than k . Let M be a maximum X_n -partition on G . Then, the class X of v is a $Star_m$, $k+1 \leq m \leq n$ with v as its center. Let $X = \{v, x_1, x_2, \dots, x_{m-1}\}$. Then, $M^* = \{M \setminus X, \{x_1\}, \{x_2\}, \dots, \{x_{m-1}\}\}$ is an X_n -partition on $G-v$. Since $m-1 \geq k > t$, M^* is not a maximum X_n -partition on $G-v$, and hence it can be augmented to a maximum Δ_n -partition, say M^{**} , such that at least t vertices in $X \setminus v$, say x_1, x_2, \dots, x_t , are exposed by M^{**} . Finally the classes in M^{**} and $\{v, x_1, x_2, \dots, x_t\}$ yield a maximum X_n -partition on G such that the class of v is a $Star_{t+1}$ centered at v , which is contradiction.

Necessity part of Statement (c): Arbitrarily take a maximum X_n -partition M on G . Let X denote the class of v under M . We will prove that the induced graph on X is a $Star_m$ with v as its center, for some m ($k+1 \leq m \leq n$). It follows from $\dim(G-v, X_n) \geq \dim(G, X_n) + k$ ($2 \leq k \leq n-1$) that $|X| \geq k+1 \geq 3$. If $(X = K_3)$ or $(X = Star_m$ ($k+1 \leq m \leq n$) and v is not the center vertex), then $\dim(G-v, X_n)$

$= \dim(G, X_n)$. Hence v must be the center vertex of $X = \text{Star}_m$, $k+1 \leq m \leq n$. ■

It should be pointed out that when $k = n-1$, this theorem reduces to Lemma 5.2.7.

Corollary 5.2.12. *Let G be a connected graph. Then,*

- (a) *If the degree of every vertex is larger than or equal to $|G|/n+1$ then G is an X_n -prime graph.*
- (b) *If the degree of every vertex is less than or equal to $n-2$, then G is an X_n -prime graph.*
- (c) *If G is an X_n -regular graph with $|G| > n$, and the degree of every vertex is less than or equal to $n-1$, then G is an X_n -prime graph.*
- (d) *$\dim(G, X_{n-1}) = 0$, then G is an X_n -prime graph.*

Proof. (a) By Theorem 5.2.9, it suffices to prove that G contains no X_n -zero. Assume, on the contrary, that y is an X_n -zero of G . Let M be a maximum X_n -partition on G . Then the class of y defined by M , say Y , is a Star_n centered at y . If one member, say x , in $Y \setminus y$ is incident to a non- Star_n -center vertex, say z , then replacing the classes Y and Z , the class of z , (by $Y \setminus \{x\}$ and $Z \cup \{x\}$, if z is the center vertex of Z) or (by $Y \setminus \{x\}$, $Z \setminus \{z\}$ and $\{x, z\}$, if z is the center vertex of Z), we have a new maximum X_n -partition, under which y is not a Star_n -center. Hence y is not an X_n -zero of G , a contradiction. Thus the vertex x is incident to only Star_n -centers. Because $\deg(x) \geq |G|/n+1$, the vertex x is incident to at least $|G|/n$ Star_n classes that are disjoint with X . Thus the size of G is at least $n(|G|/n)+n > |G|$, which is a contradiction. We then conclude that G contains no X_n -zero.

(b) From Lemma 5.2.7, that $\deg(v) \leq n-2$ for every vertex v in G implies that no vertex in G is X_n -zero. Thus by Theorem 5.2.9, G must be an X_n -prime graph.

(c) Let M be a perfect X_n -partition on G . And let v be a Star_n -center under M . Denote by $X = \{v, x_1, x_2, \dots, x_{n-1}\}$ the class of v defined by M . Because $|G| > n$, we can find a vertex, say y , which is outside X and incident to some vertex, say x_1 , in X . Since $\deg(y) \leq n-1$, y can not be a Star_n -center under M . Let Y denote the class of y under M . Then $(\{M \setminus (X \cup Y), Y \cup \{x_1\}\})$, if y is the center

vertex of its class) or $(\{M \setminus (X \cup Y), (Y \setminus y), \{x_1, y\}\}$, if y is not the center vertex of its class) is an X_n -partition on $G - v$ which exposes only $n-2$ vertices, thus

$$\dim(G - v, X_n) \leq n-2 = \dim(G, X_n) + (n-2),$$

which implies that v is not an X_n -zero. Hence there exists no X_n -zero in G , by Theorem 5.2.9, and hence G is an X_n -prime graph.

(d) Simple corollary of Lemma 5.2.7 and Theorem 5.2.9. ■

Theorem 5.2.13. *For any positive integer m , let P_m , T_m , R_m , and Z_m denote the sets of X_m -poles, X_m -infinities, X_m -roots and X_m -zeros of a given graph G with $|G|=N$, respectively. Then,*

(a)

$$T_m \subset P_{m-1}.$$

(b)

$$\emptyset = Z_N \subset R_{N-1} \subset Z_{N-1} \subset R_{N-1} \subset Z_{N-2} \subset \dots \subset R_3 \subset Z_3 \subset R_2 \subset Z_2.$$

Proof. It is obvious that $Z_N = \emptyset$, $R_2 = Z_2$ and $R_m \subset Z_m$. So it suffices to prove that $T_m \subset P_{m-1}$ and $Z_m \subset R_{m-1}$. We will only prove Statement (a), since (b) can be similarly proven. Let x be an X_m -infinity, M a maximum X_m -partition on G . Without loss of generality, we assume that M covers the vertex x . If x is an exposed vertex of G , then $x \in P_{m-1}$ follows immediately. Otherwise, starting from x , we iteratively construct a graph E whose vertices are labeled as *even vertices* and *odd vertices*.

Initially, E contains only the vertex x , which is labeled as an even vertex. We then grow E by adding vertices and edges in G .

By the definition of an X_m -infinity, the vertices adjacent to x are X_m -zeros. Let y be one of such X_m -zeros. Then the class of y under M is a Star_m centered at y , by Lemma 5.2.7. We now extend the graph E by adding the edge (x, y) and the star shaped graph induced on the whole class

of y . Among the new vertices of E , the vertex y is labeled as an *odd vertex*, while the other vertices are labeled as *even vertices*.

In general, let z be an even vertex in E . If no vertices outside E are adjacent to z , then we turn to consider the other even vertices in E . If g is a vertex of G which is adjacent to z but not a vertex of E , then the class of g defined by M is a Star_m centered at g . (If this is not the case we can find a maximum X_m -partition on G such that the class of x is no more a Star_m , which is contradictory to the assumption that x is an X_m -zero.) We then extend the graph E by adding the edge (g, z) and the star shaped graph induced on the whole class of g . Among the new vertices of E , the vertex g is labeled as an odd vertex and the other vertices are labeled even vertices.

Upon termination of this iterative process, one observes that in E , all the classes defined by M are Star_m 's. Deleting exactly one non-center vertex from each of these Star_m 's (in particular, delete vertex x) we then obtain a group of Star_{m-1} 's. Let M' be a Δ_{m-1} -partition on G with the above Star_{m-1} 's as its classes. Then through the augmenting procedure in the proof of Theorem 4.2.3, we can get a maximum X_{m-1} -partition on G which does not cover vertex x (Star_{m-1} classes are not changed in the procedure; in particular, the class of vertices adjacent to x are not changed). Thus x is an X_{m-1} -pole, and (a) then follows. ■

Theorem 5.2.14. *Let G be a connected graph and H a KB-blossom in G . Assume that G and H have the same vertex set V . Then either G is an X_3 -regular graph or G is an X_n -prime graph.*

Proof. Assume that H has more than one blocks. Let W be the vertex set of a block of H which contains a cut-vertex x of H . Let $U = (V - W) \cup \{x\}$. One checks that the induced subgraph H_U of H on U is also a KB-blossom. By induction on $|V|$, either the induced subgraph G_U of G on U possesses a perfect X_3 -partition or G_U is a complete graph. If G_U possesses a perfect X_3 -partition, we obtain an X_3 -partition on G by partitioning the set $W \setminus \{x\}$ into arbitrary pairs. So, in the

following, we assume that G_U is a complete graph.

If G is a complete graph or if it has no edges other than those of G_U and G_W (the induced subgraph of G on W), then G is a KB-blossom. The other possibility is that G has some edges between $U \setminus \{x\}$ and $W \setminus \{x\}$ but not all of them. By symmetry, let (u, w) be an edge of G and (u, w') be a non-edge, where $u \in U \setminus \{x\}$ and $w, w' \in W \setminus \{x\}$. Now, partition the set $U \setminus \{u\}$ and $W \setminus \{x, w, w'\}$ into arbitrary pairs, respectively. These pairs, together with the triple $\{u, w, w'\}$, form a perfect X_3 -partition on G . ■

References:

- [1] H. Alt, K. Mehlhorn and M. Paul, *Computing a maximum cardinality matching in a bipartite graph in time $O(n^{1.5}\sqrt{m/\log n})$* , Inform. Process. Lett., 37(1991), pp. 237-240.
- [2] A. Amahashi and M. Kano, *Factors with given components*, Discrete Math., vol. 42, pp. 1-6, 1982.
- [3] K. Ando, *Rao's conjecture on self-complementary graphs with k -factors*, J. of Graph Theory, vol. 9, pp. 119-121, 1985.
- [4] A. Akiyama and M. Kano, *Factors and factorizations of graphs, A survey*, J. of Graph Theory, vol. 9, pp. 1-41, 1985.
- [5] M. Behzad and G. Chartrand, *Introduction to the theory of graphs*, Allyn and Bacon Inc., 1971.
- [6] C. Berge, *Sur le couplage maximum d'un graphe*, C. R. Acad. Sciences, Paris, v.247,258-259, 1958.
- [7] C. Berge, *Graphs And Hypergraphs*, North-Holland Pub., 1973.
- [8] C. Berge and M. L. Vergnas, *On the existence of subgraphs with degree constraints*, Proc. Konink, Nederl. Adad. Wet. Amsterdam, 81(1978), pp. 165-176.
- [9] G. Cornu éols and W. Pulleyblank: *Perfect triangle-free 2-matchings*. Math. Prog. Study 13, 1-7 (1980).
- [10] G. Cornu éols and D. Hartvigsen, *An extension of matching theory*, J. Combin. Theory, Ser. B. 40(1986), 285-296.
- [11] G. Cornu éols, D. Hartvigsen and W. Pulleyblank, *Packing subgraphs in a graph*, Operations Research Letters (4), 139-143, 1982.

- [12] J. Edmonds, *Paths, trees and flowers*, Canad. J. Math., Vol. 17, 449-467, 1965.
- [13] J. Edmonds, *Covers and packings in a family of sets*, Bull. Amer. Math. Soc. 68, (1962), pp. 494-499.
- [14] Y. Egawa, M. Kano, A.K. Kelmans, *Star partitions of graphs*. Journal of Graph Theory, 25, 185–190 (1997).
- [15] S. Even and O. Kariv, *An $n^{2.5}$ algorithm for maximum matching in general graphs*, Proc. of the 16th Annal Sym. On Foundation of Computer Science, Berkeley, 1975, pp. 100-112.
- [16] L. Ford and D. Fulkerson, *Maximal flow through a network*, Canad. J. Math. 8(1956), 399-404.
- [17] H. N. Gabow, *A linear-time algorithm for a special case of disjoint set union*, J. of Computer and System Sciences 30. 209-221 (1985).
- [18] M. Gondran and M. Minoux, *Graphs and algorithms*, Wiley-Interscience Pub., New York, 1990.
- [19] J. Gonzalez and O. Landaeta, *A competitive strong spanning algorithm for the maximum bipartite matching problem*, SIAM J. Disc. Math., vol. 8, No, 2, pp. 186-195, 1995.
- [20] J. Graver, *Combinatorial Rigidity*, American Math. Society, 1993.
- [21] J. Graver and W. Jurkat, *f-factors and related decompositions of graphs*, J. of Combinatorial theory, Ser. B, vol. 28, pp. 66-84, 1980.
- [22] F. Harary, *Graph theory*, Addition-Welsley, 1972.
- [23] P. Hell and D. Kirkpatrick, *Star factors and star packings*, TR 82-6, Department of Computer Science, Simon Fraser University, 1982.
- [24] P. Hell and D. G. Kirkpatrick, *Packings by cliques and by finite families of graphs*, Discrete Math. 49(1984) 45-59.

- [25] P. Hell and D. G. Kirkpatrick, *On the complexity of general graph factor problems*, SIAM J. Comput. (1983), 601-609.
- [26] P. Hell and D. G. Kirkpatrick, *Packings by complete bipartite graphs*. SIAM J. Algebraic Disc. Methods. 7(1986), 113-129.
- [27] P. Hell and D. G. Kirkpatrick, *On generalized matching problems*, Inform. Process. Letters. 12(1981) 33-35.
- [28] J. E. Hopcroft and R. M. Karp, *An $n^{2.5}$ algorithm for maximum matching in bipartite graphs*, SIAM J. Comput., vol. 2, pp. 225-231, 1973.
- [29] A. K. Kelmans, *Optimal packing of induced stars in a graph*. Discrete Math. 173, 97–127 (1997).
- [30] Zoltán Király and Jácint Szabó, *Induced Graph Packing Problems*, Graphs and Combinatorics (2010) 26:243–257.
- [31] D. G. Kirkpatrick and P. Hell, *On the completeness of a generalized matching problem*, Proc. Tenth Annual ACM Symposium on Theory of Computing (1978), 240-245.
- [32] M. Las Vergnas, *An extension of Tutte's l -factor theorem*, Discrete Math. 23 (1978) 241-255.
- [33] M. Loebl and S. Poljak, *On matroids induced by packing subgraphs*, J. Combin. Theory Ser. B 44 (1988) 338-354.
- [34] M. Loebl and S. Poljak, *Subgraph packing -- A survey in: Topics in Combinatorics and Graph Theory: Essays in Honour of Gerhard Ringel* (Phisica Verlag, Wursburg, 1990) 491-503.
- [35] M. Loebl and S. Poljak, *Packing by good families of subgraphs*, Annals of Discrete Math. vol. 51, pp. 181-186, North-Holland, Amsterdam, 1992.
- [36] M. Loebl and S. Poljak, *Efficient subgraph packing*, J. Combin. Theory Ser. B 59 (1993)

- 106-121.
- [37] L. Lovasz, *Subgraphs with prescribed valencies*, J. Combin. Theory. 8(1970), 391-416.
- [38] L. Lovasz and M. Plummer, *Matching Theory*, North-Holland, Budapest, Amsterdam, 1986.
- [39] E. L. Lawler, *Combinatorial optimization: networks and matroids*, Saunders College Publishing, Chicago, 1976.
- [40] W. Lipski, Jr. and F. P. Preparata, *Efficient algorithms for finding maximum matchings in convex bipartite graphs and related problems*, Acta Informatica, 15, 329-346(1981).
- [41] C. J. Mcdiarmid, *Rado's theorem for polymatroids*, Math. Proc. Cambridge Philos. Soc. 78, 1975, 263-281.
- [42] S. Micali and V. V. Vazirani, *An $O(\sqrt{|V|}|E|)$ algorithm for finding maximum matching in general graphs*, Proc. 21st IEEE Symp. Foundations of Comput. Sci., pp. 17-27, 1980.
- [43] J. Petersen, *Die theorie der regularen graphen*, Acta Math. 15, pp. 193-220, 1891.
- [44] S.-Y. Robert Li, *The concepts of prime factorization of networks*, Proc. of Int. Workshop on Discrete Math. And algorithm (IWDMA'93), pp. 43-64; Dec. 13-14, 1993.
- [45] S.-Y. Robert Li, *Graph factorization with respect to general types of matchings*, Unpublished manuscript, Univ. of Illinois at Chicago, 1979.
- [46] S.-Y. Robert Li and Yixian Yang, *Prime factorization theory of networks*, Unpublished manuscript, The Chinese University of Hong Kong, 1992.
- [47] A. Saito and M. Watanabe, *Partitioning graphs into induced stars*. Ars Combin. 36, 3-6 (1993).
- [48] W. T. Tutte, *The factorization of linear graphs*, J. London Math. Soc., vol. 22, pp. 107-111, 1947.
- [49] W. T. Tutte, *The factors of graphs*, Canad. J. Math. 4(1952) 314-238.

- [50] V. V. Vazirani, *A theory of alternating paths and blossoms for proving correctness of the $O(\sqrt{|V|}|E|)$ general graph matching algorithm*, Private Communication.
- [51] D. J. Welsh, *Matroid Theory*, Academic Press, New York, 1976.
- [52] R. J. Wilson, *Introduction to graph theory*, Longman, New York, 1985.