GENERALIZED SATISFIABILITY WITH k OCCURRENCES PER VARIABLE: A STUDY THROUGH DELTA-MATROID PARITY

VICTOR DALMAU
UNIVERSITAT POMPEU FABRA
VICTOR.DALMAU@TECH.UPF.ES
AND
DANIEL FORD
UNIVERSITY OF CALIFORNIA
FORD@CSE.UCSC.EDU

Topic: computational complexity, generalized satisfiability, Δ-matroid

Abstract. In this paper we examine generalized satisfiability problems with variable occurrences. First, we show that 3 occurrences per variable suffice to make these problems as hard as their unrestricted version. Then we focus on generalized satisfiability problems with at most 2 occurrences per variable. It is known that some NP-complete generalized satisfiability problems become polynomially solvable when only 2 occurrences per variable are allowed. We identify two new families of generalized satisfiability problems, called local and binary, that are polynomially solvable when only 2 occurrences per variable are allowed. We achieve this result by means of a reduction to the Δ-Matroid Parity Problem, which is another important theme of this work.

1. Introduction and Summary of Results

Satisfiability problems are studied intensively in theoretical computer science. One of the main lines of research consists in identifying the restrictions of the general satisfiability problem that give rise to tractable problems. In an attempt to define a common framework to study different classes of satisfiability problems Schaefer defined the class of generalized satisfiability problems [Sch78]. He observed that many of the variants of the satisfiability problem considered in the literature, for example Horn and 2-CNF formulas, can be obtained by restricting the type of clauses that are allowed in a formula. In the generalized satisfiability framework such restrictions are implemented by means of a fixed basis, or set of Boolean relations, that define the basic building blocks from which formulas are created. Many families of satisfiability problems can be expressed using this formalism. The following example illustrates how 3-SAT can be expressed as a generalized satisfiability problem.

The 3-clause, $x \vee y \vee \neg z$, can be expressed as the atomic formula, $R(x, y, z)$, where R is a Boolean relation containing all those assignments to x, y, z satisfying $x \vee y \vee \neg z$. That is $\{0,1\}^3 \setminus \{(0,0,1)\}$. Thus, if we let the basis Γ be the set consisting of the four relations, (1) $\{0,1\}^3 \setminus \{(0,0,0)\}$, (2) $\{0,1\}^3 \setminus \{(0,0,1)\}$, (3) $\{0,1\}^3 \setminus \{(0,1,1)\}$, and (4) $\{0,1\}^3 \setminus \{(1,1,1)\}$, then $SAT(\Gamma)$ is 3-SAT. Every 3-CNF formula can be generated, clause by clause, by picking the appropriate relation from Γ and assigning to it the appropriate variables. Every finite collection of relations, Γ, defines a different generalized satisfiability problem, $SAT(\Gamma)$.

Schaefer's Dichotomy Theorem showed that generalized satisfiability problems are either in P or they are NP-complete. This is of particular interest because of Ladner's Theorem that states, if $P \neq NP$ there exists problems in NP that are neither NP-complete nor in P [Lad75]. Schaefer did even more. He completely classified the computational complexity of generalized satisfiability problems by showing that: $SAT(\Gamma)$ is in the class P if Γ is a subset of one of the following six families (1) *Horn* relations, (2) *dual-Horn* relations, (3) *bijunctive* relations, (4) *affine* relations, (5) 0-valid relations, and (6) 1-valid relations and NP-complete otherwise [Sch78].

The main motivation of this paper is to bring into the picture the number of occurrences per variable. It is well known that some NP-complete satisfiability problems become tractable when the maximum number of occurrences per variable is bounded. For example, it is easy to see that any k-CNF where each variable occurs at most k times contains at least as many variables as clauses and is consequently satisfiable. This result has been extended in [KST93]. In particular Kratochvíl et al. showed that for every k there exists an exponential function, $l(k) \geq k$, such that k-SAT with $\leq l(k)$ variable occurrences, and at most one occurrence per clause, is *trivial*, that is, every instance is satisfiable, but k-SAT with $\leq l(k)+1$ occurrences per variable is NP-complete.

The fact that some NP-complete satisfiability problems become solvable in polynomial time when the number of occurrences per variable is bounded and the fact that the generalized satisfiability framework is a well established setting for the uniform analysis of satisfiability problems motivated the study of generalized satisfiability problems with a bounded number of occurrences per variable, introduced independently in [Fed01] and [Ist97]. The ultimate goal here is to find a refinement of Schaefer's Dichotomy Theorem in which the maximum number of occurrences per variable is taken into account. More precisely, if we let $SAT(k, \Gamma)$ contain those formulas (instances) of $SAT(\Gamma)$ in which every variable occurs at most k times, then we are interested in determining, for every positive integer k and every basis Γ, the computational complexity of $SAT(k, \Gamma)$.

Matching on graphs with bounded degree is a typical example of a problem that can be expressed as a generalized satisfiability problem with a bounded number of occurrences per variable: Edges of the graph are represented by variables. Each vertex of degree j is represented by an ordered set of j variables, one per incident edge, and the j-ary Boolean relation *1-in-j*: the j-element set of all j-dimensional unit vectors, e.g. the relation 1-in-3 is $\{(1,0,0), (0,1,0), (0,0,1)\}$.

A 1 in the relation corresponds to an edge that is assigned to the final matching and a 0 corresponds to an edge that is not. The relation 1-in-j constrains variable assignments to those that assign exactly one incident edge, of the corresponding vertex, to the final matching. Since every edge is incident to two vertices every variable must occur exactly twice. If we let Γ_m be the set consisting of the m relations 1-in-j for $1 \leq j \leq m$, then $SAT(2, \Gamma_m)$ can express every instance of the matching problem on graphs with degree $\leq m$. This is a subclass of the graph matching problem which Edmonds' Matching Algorithm solves in polynomial time [Edm65]. However, Schaefer showed that the relation 1-in-3 is neither Horn, dual-Horn, bijunctive, affine, 0-valid, nor 1-valid, thereby proving that $SAT(\Gamma_m)$ is NP-complete whenever $m \geq 3$. [Sch78]. Edmonds was able to show tractability for the graph matching problem whereas Schaefer's Theorem does not because Edmonds' Matching Algorithm is concerned with graphs, as opposed to hypergraphs, where each edge is incident with exactly two vertices. This corresponds, in the generalized satisfiability framework, to each variable occurring exactly twice. In other words, although it is possible to express graph matching problems with bounded degree as generalized satisfiability problems without restricting the number of occurrences, the polynomial solvability of this problem cannot be explained by means of this reduction. However, if we bound the number of occurrences to 2 then the situation turns out to be very different. In particular, Istrate [Ist97] first noticed that $SAT(2, \Gamma_m)$ reduces in polynomial time to the matching problem on graphs of degree at most m. Consequently, $SAT(2, \Gamma_m)$ captures perfectly the computational complexity of graph matching problems on graphs of degree at most m.

Graph matching is not an isolated example of a generalized satisfiability problem that is tractable when variables occur at most twice and NP-complete otherwise. In particular, it is known that if every relation of a basis, Γ, is compact then $SAT(2, \Gamma)$ is solvable in polynomial-time whereas $SAT(\Gamma)$ is, in general, NP-complete [Ist97]. Another example of a broad family of generalized satisfiability problems that becomes tractable when the number of occurrences is bounded by 2 is the class of bases consisting only of co-independent relations [Fed01]. Generalized satisfiability with at most two occurrences per variable is a rich class of problems intimately related to algorithms, concepts and methods of matching theory. Moreover, the computational complexity of these problems remains largely unknown.

In order to simplify matters we only consider in this paper a variant of the generalized satisfiability problems, denoted $SAT_C(\Gamma)$, already introduced by Schaeffer, in which constants 0 and 1 can be used in the formulas. This assumption, widely used in the literature (see [CKS01]) allows us to simplify considerably the picture. In particular, we are able to show that if $k \geq 3$ then $SAT_C(k, \Gamma)$, is polynomially equivalent to $SAT_C(\Gamma)$ for every basis Γ.

This result shows that three occurrences suffice to make a problem as hard as the unrestricted problem. As an immediate corollary we get that many common satisfiability problems such as 1-in-3 SAT and not all equal 3-SAT (see [GJ79] for definitions), are NP-complete with ≤ 3 occurrences per variable (these problems are known to be tractable with ≤ 2 occurrences).

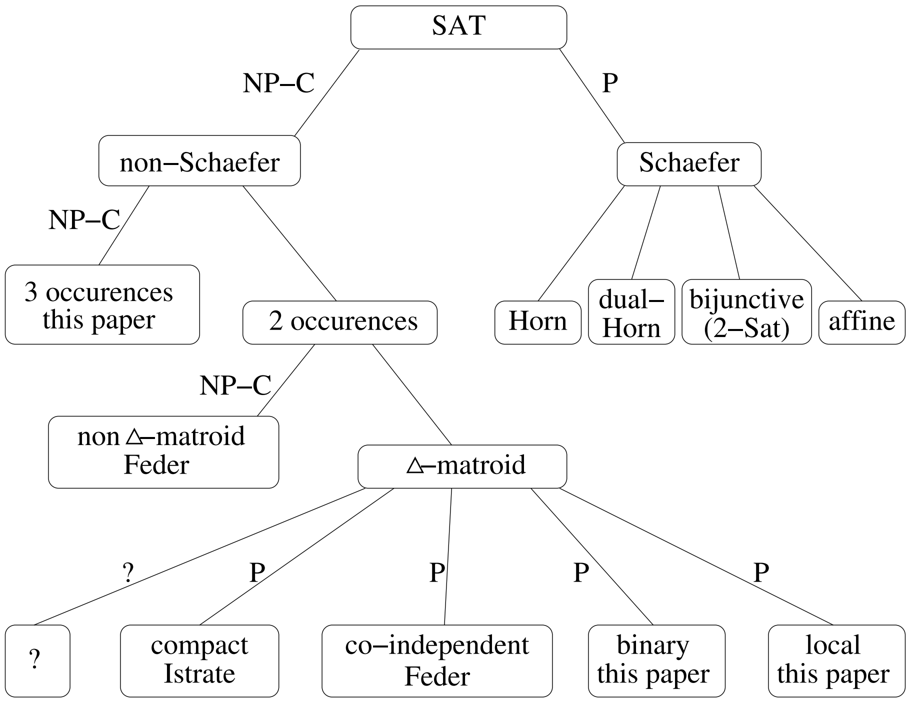

FIGURE 1.1. Computational complexity of generalized satisfiability with limited occurrences per variable

Since generalized satisfiability problems with exactly one occurrence per variable are trivially solvable because every instance is satisfiable, we only need to solve the case when $k = 2$. It was recently shown that $SAT_C(2,\Gamma)$ is polynomially equivalent to $SAT_C(\Gamma)$ if Γ contains a relation that is not a Δ-matroid [Fed01]. These two results combine to identify $SAT_C(2,\Gamma)$ where every relation in Γ is a Δ-matroid as the only family that may contain additional tractable families that were not previously identified by Schaefer. Thus, at this point, the only task that remains to be done is to identify all the bases, Γ, containing only Δ-matroids such that $SAT_C(2,\Gamma)$ is tractable. As we mentioned above, only two families of such basis have been identified so far (prior to the results in this paper): the class of bases containing only compact relations and the class of bases containing only of co-independent relations [Ist97, Fed01]. In this paper we add two new families of bases to the list: the class of bases containing only *local* Δ*-matroids* and the class of basis containing only *binary* Δ*-matroids*. The strategy used to obtain these results consists in a reduction to the Δ-matroid parity problem (see [GIM97] for definitions), so that tractable cases of the latter could be transformed to tractable cases of $SAT_C(2,\Gamma)$. The only previously known tractable cases of the Δ-matroid parity problem are, up to our knowledge, *binary* and *linear* Δ-matroids [LP86, GIM97]. However, only one of these two tractable cases, namely the family of binary Δ-matroids, can be moved into a tractable case of $SAT_C(2,\Gamma)$. The reason for that is that all the known algorithms for these classes depend on the fact that a representation, in terms of matrices, of the Δ-matroid is known. Whereas for binary Δ-matroids such a matrix representation can be constructed, it is still open whether this is possible for linear Δ-matroids. On the other hand, the class of local Δ-matroids constitutes a brand new family of tractable cases of the Δ-matroid parity problem. The proof of this result relies in a generalization of the concept of augmenting paths originally introduced by Berge in graph matching [Ber57].

Figure 1.1 illustrates the current state of knowledge about the computational complexity of $SAT_C(2,\Gamma)$.

Because of space restrictions most proofs have been moved to the appendix, where they can be read at the referee's discretion.

2. Generalized satisfiability with at most k occurrences per variable

We introduced generalized satisfiability problems, $SAT(\Gamma)$, via the example of 3-SAT where clauses were replaced by an ordered set of variables and a relation. We now give a formal definition of $SAT(\Gamma)$. An *r-ary relation R* is any nonempty subset of $\{0,1\}^r$, R' is a *predicate symbol* with the same arity as R, and a *basis* or *constraint set* Γ is a finite collection $\{R_1, \ldots, R_m\}$ of relations. A *$CNF(\Gamma)$-formula* is a finite conjunction of clauses $C_1 \wedge \cdots \wedge C_n$ such that each clause, C_i, is an *atomic formula* of the form $R'(v_1, \ldots, v_r)$ where $v_1, \ldots, v_r$ are Boolean variables in an infinite set V, and R is an r-ary relation in Γ. An atomic formula $R'(v_1, ..., v_r)$ is *satisfied* by a variable assignment $f : V \to \{0,1\}$ if and only if $(f(v_1), ..., f(v_r)) \in R$, and a $CNF(\Gamma)$-formula is *satisfiable* if and only if there exists an assignment satisfying all its clauses. It is sometimes customary to assume that we can replace some variables in a $CNF(\Gamma)$-formula by the constant symbols $\mathbf{0}$ and $\mathbf{1}$ to be interpreted as 0 and 1 respectively. We call any formula obtained this way a $CNF_C(\Gamma)$-formula. Each basis Γ gives rise to the *generalized satisfiability problem* $SAT(\Gamma)$: given a $CNF(\Gamma)$-formula, is it satisfiable? The problem *generalized satisfiability problem with constants*, $SAT_C(\Gamma)$, is defined similarly.

As we mentioned in the introduction, Schaefer completely classified the computational complexity of generalized satisfiability problems [Sch78]. In order to state Schaefer's Dichotomy Theorem (Theorem 2.2 below) we introduce the following definitions, where $\vee, \wedge$, and $\oplus$ act on tuples component-wise.

Definition 2.1. A relation, R, is *Horn* if $\mathbf{x}, \mathbf{y} \in R \Rightarrow \mathbf{x} \wedge \mathbf{y} \in R$, *dual-Horn* if $\mathbf{x}, \mathbf{y} \in R \Rightarrow \mathbf{x} \vee \mathbf{y} \in R$, *bijunctive* if $\mathbf{x}, \mathbf{y}, \mathbf{z} \in R \Rightarrow (\mathbf{x} \wedge \mathbf{y}) \vee (\mathbf{x} \wedge \mathbf{z}) \vee (\mathbf{y} \wedge \mathbf{z}) \in R$, *affine* if $\mathbf{x}, \mathbf{y}, \mathbf{z} \in R \Rightarrow \mathbf{x} \oplus \mathbf{y} \oplus \mathbf{z} \in R$, 1-*valid* if it contains the tuple $(1, 1, \ldots, 1)$, and 0-*valid* if it contains the tuple $(0, 0, \ldots, 0)$. A basis, Γ, is Horn (respectively, dual-Horn, bijunctive, affine, 1-valid, or 0-valid) if every relation in Γ is Horn (respectively, dual-Horn, bijunctive, affine, 1-valid, or 0-valid). We say that Γ is *Schaefer* if it is Horn, dual-Horn, bijunctive or affine.

Theorem 2.2. [Sch78] *Let Γ be a finite set of logical relations.*
If Γ is Schaefer, 1*-valid, or* 0*-valid then $SAT(\Gamma)$ is in P; otherwise, it is NP-complete.*
If Γ is Schaefer, then $SAT_C(\Gamma)$ is in P; otherwise, it is NP-complete.

The ultimate goal of this paper is to completely classify the complexity of generalized satisfiability problems (with and without constants) when the number of occurrences per variable is bounded. Although we did not completely achieve this objective we have been able to identify some subclasses of the NP-complete problems identified by Schaefer that become tractable when the number of occurrences per variables is bounded. In what follows we will only deal with generalized satisfiability problems with constants, $SAT_C(\Gamma)$. The first question we asked is how many occurrences per variable suffice to make the restricted problem as hard as the general problem. The answer, we found, is 3.

2.1. 3 occurrences per variable can simulate any number of occurrences.

Let $k \geq 1$ be a positive integer. We define *$CNF_C(k,\Gamma)$-formulas* to be the subset of $CNF_C(\Gamma)$-formulas restricted to those formulas where each variable occurs at most k times. $SAT_C(k,\Gamma)$ is defined to be the following decision problem: given a $CNF_C(k,\Gamma)$-formula ϕ, is it satisfiable? $SAT(k,\Gamma)$ is defined similarly.

The following theorem could be derived from Theorem 3 in [Fed01]. We include a self-contained proof in the appendix for the sake of completeness.

Theorem 2.3. *Let Γ be a finite set of logical relations, and let $k \geq 3$. Then $SAT_C(k,\Gamma)$ is polynomially equivalent to $SAT_C(\Gamma)$.*

Proof. See Appendix A. □

An immediate consequence of this theorem is that, when constants are allowed, three occurrences suffice for Positive 1 in 3 SAT and Not all Equal SAT (refer to [GJ79] for definitions) to be NP-complete.

In view of the results of Kratochvíl et al. mentioned in the introduction, the state of affairs for generalized satisfiability without constants is more complicated [KST93].

2.2. At most 2 occurrences per variable.

Any generalized satisfiability problem with at most one occurrence per variable, $SAT_C(1,\Gamma)$ for any basis Γ, is trivially solvable in polynomial time, that is, every instance is satisfiable. Therefore, the only case of unknown complexity is when variables occur at most twice.

The computational complexity of $SAT_C(2,\Gamma)$ appears to be quite interesting. Let $\{$1-in-3$\}$ be the basis containing the single relation 1-in-3 given by $\{(0,0,1),(0,1,0),(1,0,0)\}$. Istrate first observed that $SAT_C(2,\{\text{1-in-3}\})$ is solvable in polynomial-time, by means of a reduction to the matching problem in graphs, whereas $SAT_C(\{\text{1-in-3}\})$ is NP-complete [Sch78, Ist97]. The guiding purpose of this work is to identify those bases, Γ, like $\{$1-in-3$\}$, such that $SAT_C(2,\Gamma)$ is polynomially solvable and $SAT_C(\Gamma)$ is NP-complete. That is, we ask, *assuming $P \neq NP$, when is tractability contingent on having at most two occurrences per variable?*

2.2.1. *Previous results.* Δ*-matroids.* When variables occur at most twice there are two previously known tractable classes in addition to those identified by Schaefer. $SAT_C(2,\Gamma)$ is polynomially solvable when every relation in Γ is *compact* or when every relation in Γ is *co-independent* [Ist97, Fed01]. An n-ary relation, R, is co-independent if $d(\mathbf{x},\mathbf{y}) \neq 1$ for all $\mathbf{x},\mathbf{y} \in \{0,1\}^n \setminus R$ where d is the hamming distance. We do not provide a definition of compact relations since it is not needed to obtain our results and it is rather involved. On the other hand, it is also known that if a basis, Γ, contains a relation which is not a Δ*-matroid* relation, then $SAT_C(2,\Gamma)$ is not any easier that $SAT_C(\Gamma)$ [Fed01]. Thus, any new tractable basis, that is, any non-Schaefer basis, Γ, such that $SAT_C(2,\Gamma)$ is solvable in polynomial time, must consists entirely of Δ-matroids. We now define Δ-matroids which will play a central role in our study.

Definition 2.4. [BC95] Let E be a finite set and $\mathcal{F} \subseteq \mathcal{P}(E)$ a collection of subsets of E. The pair $(E,\mathcal{F})$ is called a *set system*. The set E is called the *universe* and the subsets in $\mathcal{F}$ are called the *feasible* sets of $(E,\mathcal{F})$. Let Δ be the *symmetric difference operator*, that is, for two sets A and B, $A\Delta B = (A\cup B)\setminus(A\cap B)$. Then $(E,\mathcal{F})$ is a Δ*-matroid* if $\mathcal{F}$ satisfies the following *symmetric exchange axiom*:

$\forall A,B \in \mathcal{F}$ and $\forall x \in A\Delta B$, $\exists y \in A\Delta B$ such that $A\Delta\{x,y\} \in \mathcal{F}$. Notice that y is not necessarily different from x.

Δ-matroids are a well studied generalization of matroids [BJ00, Wen99, BC95, Bou89]. Let us introduce here a running example that illustrates how Δ-matroids can be used to express matching problems.

Example 2.5. Let $G = (V,H)$ be an undirected graph (See Fig. 3.1.A). We will construct from G a set system $(E,\mathcal{F})$ that will turn out to also be a Δ-matroid. The set E is given by the set of pairs $(v,e) \in V\times H$ such that $v \in e$. Intuitively, we can regard E as the set of edges that we obtain when we construct a new graph G' from G by splicing up every edge in H into two edges that share a common new node (See Fig. 3.1.B). The set $\mathcal{F}$ is defined to contain those $F \subseteq E$ such that for every $v \in V$, $|F \cap (\{v\}\times H)| = 1$. Intuitively, we can regard the feasible sets in $\mathcal{F}$ as those sets of edges of G' that satisfy the matching condition over the nodes in G (See Fig. 3.1.C and 3.1.E). It is sraightforward to verify that $(E,\mathcal{F})$ satisfies the symmetric exchange axiom.

Δ-matroids can also be viewed in terms of relations by associating feasible sets with their indicator tuples.

Definition 2.6. [Lov97] Let $R \subseteq \{0,1\}^r$ be a relation. Let $\mathbf{x},\mathbf{y},\mathbf{x}' \in \{0,1\}^r$, then $\mathbf{x}'$ is a *step from* $\mathbf{x}$ *to* $\mathbf{y}$ if $d(\mathbf{x},\mathbf{x}') = 1$ and $d(\mathbf{x},\mathbf{x}') + d(\mathbf{x}',\mathbf{y}) = d(\mathbf{x},\mathbf{y})$. R is a Δ-matroid (relation) if it satisfies the following *two-step axiom:*

$\forall \mathbf{x},\mathbf{y} \in R$ and $\forall \mathbf{x}'$ a step from $\mathbf{x}$ to $\mathbf{y}$, either $\mathbf{x}' \in R$ or $\exists \mathbf{x}'' \in R$ which is a step from $\mathbf{x}'$ to $\mathbf{y}$.

A basis, Γ, will be called a Δ*-matroid basis* if every relation in Γ is a Δ-matroid relation.

Notice that a Δ-matroid relation can be obtained as the set of indicator tuples of a Δ-matroid. More formally, let $(E,\mathcal{F})$ be a Δ-matroid where $E = \{u_1,\ldots,u_n\}$. Thus, the n-ary relation R containing for every $F \in \mathcal{F}$ the tuple $\mathbf{t}_F$ given by $\mathbf{t}_F[i] = 1$ if $u_i \in F$ and 0 otherwise is a Δ-matroid relation. We say that R is obtained from $(E,\mathcal{F})$ (via the ordering $u_1,\ldots,u_n$). Conversely, given a n-ary relation R we can construct a Δ-matroid $(E,\mathcal{F})$ in the following way: $E = \{1,\ldots,n\}$ and $\mathcal{F}$ contains a feasible set F if and only if there exists a tuple $\mathbf{t} \in R$ such that $F = \{1 \leq i \leq n \mid \mathbf{t}[i] = 1\}$. Consequently, Δ-matroids and Δ-matroid relations are essentially different mathematical embodiments of the same concept.

We will change freely between the set system and the relation formalism of Δ-matroids as convenient. In the process we will often abuse notation and let R denote both the Δ-matroid relation and the Δ-matroid from which it is obtained.

Now we have introduced all the necessary conceptual machinery to state formally Feder's result.

Theorem 2.7. [Fed01] *Let Γ be a non-Δ-matroid basis. Then $SAT_C(2,\Gamma)$ is polynomially equivalent to $SAT_C(\Gamma)$.*

3. The Δ-matroid parity problem

When Γ is a Δ-matroid basis the problem $SAT_C(2,\Gamma)$ is intimately related with a problem in combinatorics called the Δ-matroid parity problem. The Δ-matroid parity problem was shown to be equivalent to the Δ-covering problem and is a generalization of the well known matroid parity problem [GIM97, Bou95, LP86].

Definition 3.1. Let $(E,\mathcal{F})$ be a Δ-matroid and $\mathcal{L}$ a partition of E into pairs. For every $u \in E$, its *mate* will be denoted by $\overline{u}$, that is, $\overline{u}$ is the only element in E such that $\{u,\overline{u}\} \in \mathcal{L}$. Let $F \in \mathcal{F}$ be a feasible set. We will let $\mathcal{L}_F$ denote the subset of $\mathcal{L}$ containing those pairs $\{u,\overline{u}\} \in \mathcal{L}$ such that either both u and $\overline{u}$ are in F or neither u nor $\overline{u}$ are in F. An instance of the Δ-matroid parity problem consist of a Δ-matroid, $(E,\mathcal{F})$, and a partition, $\mathcal{L}$, of E into pairs. The goal is to find a feasible set $F \in \mathcal{F}$ such that $\mathcal{L}_F$ is maximum, that is, at least as large as $\mathcal{L}_G$ for any other $G \in \mathcal{F}$.

Example 3.2. Let $G=(V,H)$ and $(E,\mathcal{F})$ as defined in Example 2.5. Let us define $\mathcal{L}$ as the collection of all pairs $\{(v_1,e),(v_2,e)\} \in E$, where $\{v_1,v_2\}=e \in H$. So that $|\mathcal{L}|=|H|$ and $\mathcal{L}$ is a partition of E into pairs. Thus, given a feasible set $F \in \mathcal{F}$, $\mathcal{L}_F$ denotes the collection of edges $e=\{v_1,v_2\}$ in H such that either both (v_1,e) and (v_2,e) are in F or neither is in F. Note, $|\mathcal{L}_F| \leq |\mathcal{L}| = |H|$.

Given a complete matching, $J \subseteq H$ of G, we can construct a feasible set, F, such that $\mathcal{L}_F=\mathcal{L}$. This feasible set contains those $(v,e) \in E$ such that $e \in J$. Conversely, given any feasible set F with $\mathcal{L}_F=\mathcal{L}$, the set $J \subseteq H$ containing those e such that $(v,e) \in F$ for some v, defines a complete matching of G (See Fig. 3.1.E and 3.1.F).

For computational complexity purposes, the Δ-matroid parity problem, as defined above, is not adequately described, since it is not clear what the input of the problem is, that is, *how is the Δ-matroid given?* If the Δ-matroid is specified in its entirety then the problem is trivially tractable by exhaustive search. At the other end of the spectrum a Δ-matroid, $(E,\mathcal{F})$, can be specified via a *feasible oracle*. That is, given $F \subseteq E$ the oracle tells whether F is in $\mathcal{F}$ or not. Unless explicitly stated we will assume that the Δ-matroid is specified by means of a feasible oracle. We will also assume that one feasible set is always available. Using this representation, we say that the Δ-matroid parity problem is polynomially solvable if it can be solved in time polynomial in the size $|E|$ of the universe of the Δ-matroid. Most of the work on the Δ-matroid parity problem (and on the matroid parity problem) uses this representation (see for example [GIM97]).

Lovász showed that the Δ-matroid parity problem, with feasible oracle representation, requires time exponential in the size of the universe even if the Δ-matroid of the instance is restricted to be a matroid [Lov81]. However, the problem is polynomial-time solvable if the Δ-matroid in an instance is either *linear* or *binary* and is specified by its matrix representation [GIM97]). In order to make precise the intimate relationship between generalized satisfiability problems with at most two occurrences per variables and the Δ-matroid parity problem we introduce the following definitions.

Definition 3.3. Let $\mathcal{S}$ be a collection of Δ-matroids. The $\mathcal{S}$-parity problem contains all the instances $((E,\mathcal{F}),\mathcal{L})$ of the Δ-matroid parity problem such that $(E,\mathcal{F})$ belongs to $\mathcal{S}$. Let $(E_1,\mathcal{F}_1)$ and $(E_2,\mathcal{F}_2)$ be two Δ-matroids such that $E_1 \cap E_2 = \emptyset$. Let $E=E_1 \cup E_2$ and $\mathcal{F}=\{F_1 \cup F_2 : F_1 \in \mathcal{F}_1, F_2 \in \mathcal{F}_2\}$. Then $(E,\mathcal{F})$ is the *direct sum* of $(E_1,\mathcal{F}_1)$ and $(E_2,\mathcal{F}_2)$. It is easy to see that the direct sum of two Δ-matroids is a Δ-matroid. A set $\mathcal{S}$ of Δ-matroids is *closed under direct sum* if the direct sum of any two elements in $\mathcal{S}$ is an element in $\mathcal{S}$. Two Δ-matroids $(E_1,\mathcal{F}_1),(E_2,\mathcal{F}_2)$ are called *isomorphic* if there exists some bijection $h: E_1 \to E_2$ such that for every $u_1,\ldots,u_n \in E_1$, $\{u_1,\ldots,u_n\} \in \mathcal{F}_1$ iff $\{h(u_1),\ldots,h(u_n)\} \in \mathcal{F}_2$. A set $\mathcal{S}$ of Δ-matroids is *closed under isomorphism* if every Δ-matroid $(E_2,\mathcal{F}_2)$ isomorphic to an element $(E_1,\mathcal{F}_1)$ of $\mathcal{S}$ belongs also to $\mathcal{S}$.

Theorem 3.4. *Let $\mathcal{S}$ be any collection of Δ-matroids closed under direct sum and under isomorphism, such that the $\mathcal{S}$-parity problem is polynomially solvable. Let Γ be a basis containing relations only from $\mathcal{S}$. Then $SAT_C(2,\Gamma)$ is in P.*

Proof. See Appendix B. □

All new tractable cases of $SAT_C(2,\Gamma)$ introduced in this paper, namely local and binary are obtained by identifying some polynomial-time solvable cases of the $\mathcal{S}$-parity problem and transforming them into tractable cases of the generalized satisfiability problem with at most 2 occurrences per variable.

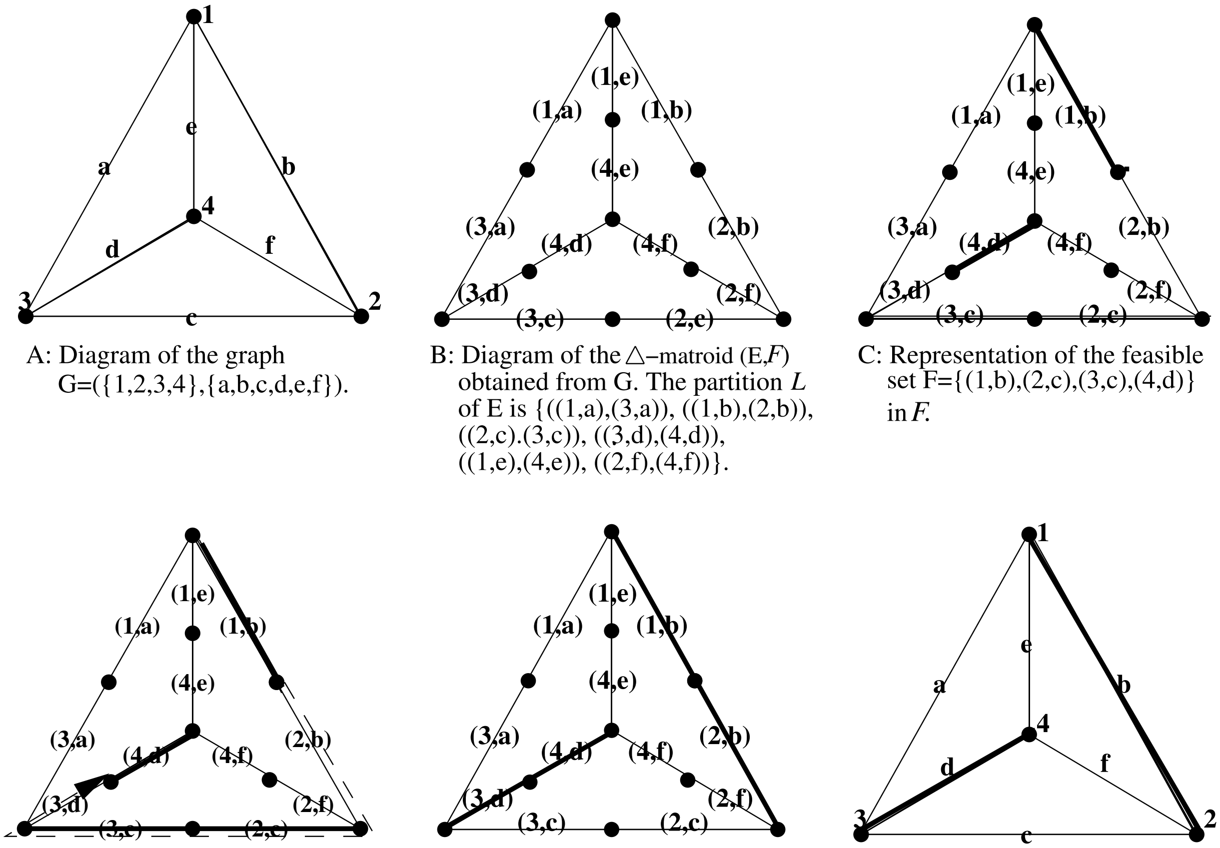

A: Diagram of the graph G=({1,2,3,4},{a,b,c,d,e,f}).

B: Diagram of the △–matroid (E,F) obtained from G. The partition L of E is {((1,a),(3,a)), ((1,b),(2,b)), ((2,c).(3,c)), ((3,d),(4,d)), ((1,e),(4,e)), ((2,f),(4,f))}.

C: Representation of the feasible set F={(1,b),(2,c),(3,c),(4,d)} in F.

D: P=(2,b),(2,c),(3,c),(3,d) is an F–augmenting L_F–alternating path.

E: Representation of the completely matched feasible set F'={(1,b),(2,b),(3,d),(4,d)}= F△ {(2,b),(2,c),(3,c),(3,d)} in F

F: Matching of G that corresponds to F'.

FIGURE 3.1. Augmenting Path of a Δ-matroid

3.1. **The Method of Augmenting Paths.** Beginning with Edmonds' original matching algorithm (see [Edm65]), the notion of an augmenting path has been central to most matching algorithms. In this section we generalize the notion of an augmenting path in order to make it suitable for the Δ-matroid parity problem.

Definition 3.5. Let $((E,\mathcal{F}),\mathcal{L})$ be an instance of the Δ-matroid parity problem. A *path* in $(E,\mathcal{F})$ (or simply a path when the Δ-matroid is implicit) is an ordered collection, $u_1,\ldots,u_n$, of different elements in E. Let $L \subseteq \mathcal{L}$ be any collection of pairs of $\mathcal{L}$. A path $u_1,\ldots,u_n$ is called *L-alternating* if: (1) for every $1 \leq 2j < n$, $\{u_{2j},u_{2j+1}\} \in L$, (2) $\{u_1,\overline{u_1}\} \notin L$, and (3) if n is even then $\{u_1,u_n\},\{u_n,\overline{u_n}\} \notin L$. Let $F \in \mathcal{F}$ be any feasible set. We say that a path $u_1,\ldots,u_n$ is an *F-augmenting path* (or simply an augmenting path when F is implicit) if: (1) $F\Delta\{u_1,\ldots,u_{2j}\} \in \mathcal{F}$ for all $1 < 2j \leq n$ and (2) $F\Delta\{u_1,\ldots,u_n\} \in \mathcal{F}$.

The basic intuition behind this definition is that if F is a feasible set such that $|\mathcal{L}_F|$ is not maximum, then there exists some F-augmenting $\mathcal{L}_F$-alternating path $u_1,\ldots,u_n$. This path can be used to obtain a new feasible set $G = F\Delta\{u_1,\ldots,u_n\}$ which increases the objective function that we intend to maximize ($|\mathcal{L}_G| > |\mathcal{L}_F|$) (See Fig. 3.1.C, 3.1.D, and 3.1.E).

Theorem 3.6. *Let $((E,\mathcal{F}),\mathcal{L})$ be an instance of the Δ-matroid parity problem and $F \in \mathcal{F}$ such that $|\mathcal{L}_F|$ is not maximum. Then there exists an F-augmenting $\mathcal{L}_F$-alternating path. Furthermore, we can compute it in time polynomial in $|E|$ given a $G \in \mathcal{F}$ with $|\mathcal{L}_G| > |\mathcal{L}_F|$.*

Proof. See Appendix C. $\square$

Let $\mathcal{S}$ be a collection of Δ-matroids. The *$\mathcal{S}$-augmenting path problem* is: given a Δ-matroid $(E, \mathcal{F})$ in $\mathcal{S}$, a partition $\mathcal{L}$ of E into pairs, and a feasible set $F \in \mathcal{F}$, find an F-augmenting $\mathcal{L}_F$-alternating path or determine that none exists.

Just like with the Δ-matroid parity problem, the computational complexity of this problem is ill defined until we specify exactly how the Δ-matroid is given. Similarly with Δ-matroid parity problems we will assume, unless otherwise explicitly stated, that the Δ-matroid is given as a feasible oracle.

Corollary 3.7. *For every collection $\mathcal{S}$ of Δ-matroids, the $\mathcal{S}$-augmenting path problem is polynomially equivalent to the $\mathcal{S}$-parity problem.*

Proof. See Appendix C. □

Recall that Lovász showed that the matroid parity problem with feasible oracle requires time exponential in the arity of the matroid [Lov81]. Hence, the matroid augmenting path problem with feasible oracle also requires time exponential in the arity of the matroid.

3.2. **A new tractable case of Δ-matroid parity: local Δ-matroids.** As we have seen, the $\mathcal{S}$-parity problem—and hence the $SAT_C(2, \Gamma)$ problem when $\Gamma \subseteq \mathcal{S}$—reduces in polynomial time to the problem of either finding an augmenting path or determining that none exists. Although, in general, this problem requires exponential time, for some sets, $\mathcal{S}$, it can be done in polynomial time. In this section we show that if a Δ-matroid satisfies an additional property, i.e. it is local, then the existence of an augmenting path is equivalent to the existence of a specific type of path, called a local path. Furthermore, we show that it takes only polynomial time to (1) find an alternating local path or determining that none exists and (2) transform an alternating local path into an alternating augmenting path. Consequently, we conclude that the $\mathcal{S}$-parity problem is tractable when $\mathcal{S}$ consists of local Δ-matroids.

Definition 3.8. Let $(E, \mathcal{F})$ be a Δ-matroid and let $F \in \mathcal{F}$ be a feasible set. A path $u_1, \ldots, u_n$ is called an *F-local path* (or simply a local path when F is implicit) if: (1) $F \Delta \{u_{2j-1}, u_{2j}\} \in \mathcal{F}$ for all $1 < 2j \leq n$ and (2) $F \Delta \{u_n\} \in \mathcal{F}$ if n is odd. A *cut* of $u_1, \ldots, u_n$ is any path $u_1, \ldots, u_j$ where $1 \leq j \leq n$, and j is odd. A *shortcut* of $u_1, \ldots, u_n$ is any path $u_1, \ldots, u_j, u_k, u_{k+1}, \ldots, u_n$ where $1 \leq j < k \leq n$, j is odd, and k is even. A *subpath* of $u_1, \ldots, u_n$ is any path obtained from $u_1, \ldots, u_n$ by a (possibly empty) sequence of cuts and shortcuts. A path is *F-local-minimal* if it is F-local and it does not contain any proper F-local subpath. Let $(E, \mathcal{F})$ be a Δ-matroid relation, and $M \geq 0$. $(E, \mathcal{F})$ is called *M-local* if for every $F \in \mathcal{F}$, every F-local-minimal path of length at most $2M$ is also F-augmenting. $(E, \mathcal{F})$ is called *local* if it is M-local for every $M \geq 0$.

Example 3.9. It is not difficult to see that the Δ-matroid $(E, \mathcal{F})$ defined in Example 3.2 is local.

Let $\mathcal{S}L$ be the set of all local Δ-matroids.

Theorem 3.10. *The $\mathcal{S}L$-parity problem is polynomially solvable.*

Proof. See Appendix D. □

Furthermore, as it is shown in Appendix E, $\mathcal{S}L$ is closed under direct sum. Thus, we obtain the following corollary. A Δ-matroid basis Γ will be called a *local basis* if every relation in Γ is local.

Corollary 3.11. *$SAT_C(2, \Gamma)$ is in P* if *Γ is local.*

3.3. **A new tractable case of generalized satisfiability with at most 2 occurrences per variable: binary Δ-matroids.** As we mentioned in Section 3 binary and linear Δ-matroids are also polynomially solvable cases of the Δ-matroid parity problem.

As we did with local Δ-matroids, we would like to use Theorem 3.4 to transform the tractability of the linear (binary) Δ-matroid parity problem into corresponding additional tractable families of $SAT_C(2, \Gamma)$. Unfortunately, Theorem 3.4 doesn't apply in these cases since the linear (binary) Δ-matroid parity problems are only known to be tractable when the instances include a matrix representation of the Δ-matroid. However we can work around this problem, in the binary case, and get an additional tractable family of $SAT_C(2, \Gamma)$ when every element in Γ is a binary Δ-matroid. We continue this section with the definition of *the linear (binary) Δ-matroid parity problem*, and conclude with Theorem 3.15, which provides the desired additional tractable case of $SAT_C(2, \Gamma)$.

Definition 3.12. Let $R \subseteq \{0,1\}^n$ be a Δ-matroid and $\mathbf{x} \in \{0,1\}^n$. Then $R \oplus \mathbf{x} := \{\mathbf{x} \oplus \mathbf{y} : \mathbf{y} \in R\}$ is the *twisting* of R by $\mathbf{x}$. It is easy to see that $R \oplus \mathbf{x}$ is a Δ-matroid. Two Δ-matroids are said to be *equivalent* if one is a twisting of the other. A matrix is said to be *skew-symmetric* if all its diagonal entries are 0 and $A = -A^T$. A matrix is said to be *symmetric* if $A = A^T$. Hence, skew-symmetric matrices over $GF(2)$ are a proper subset of symmetric matrices over $GF(2)$. Let A be a (skew-)symmetric n by n matrix. For $\mathbf{x} \in \{0,1\}^n$, define $A[\mathbf{x}]$ to be the principal sub-matrix of A indexed by the non-zero elements of $\mathbf{x}$. Define $\mathcal{R}(A) := \{\mathbf{x} \in \{0,1\}^n : \text{rank } A[\mathbf{x}] = \|\mathbf{x}\|_1\}$. Then $\mathcal{R}(A)$ is an n-ary Δ-matroid relation [Bou88].

A Δ-matroid R is called *linear* if it is equivalent to $\mathcal{R}(A)$ for some skew-symmetric matrix A, and *binary* if it is equivalent to $\mathcal{R}(A)$ for some symmetric matrix A over $GF(2)$ [GIM97]. $(A, \mathbf{x})$ will be called the *matrix representation* of the linear or binary Δ-matroid $\mathcal{R}(A) \oplus \mathbf{x}$ (if $(A, \mathbf{x})$ is the *matrix representation* of R then $\mathbf{x} \in R$.)

Definition 3.13. The *linear (binary) Δ-matroid parity problem* is: given a matrix representation for a linear (binary) n-ary Δ-matroid, R, and, $\mathcal{L}$, a partition of $\{1, \ldots, n\}$ into pairs, find the maximum matching in R with respect to $\mathcal{L}$.

Theorem 3.14. [GIM97] *The linear and binary Δ-matroid parity problems with matrix representations are in P.*

From which we obtain the following corollary. A Δ-matroid basis Γ will be called *binary* if every relation in Γ is a binary Δ-matroid.

Theorem 3.15. $SAT_C(2, \Gamma)$ *is in P* if Γ *is binary.*

Proof. See Appendix B. ☐

4. Forbidden Minors

The method of characterizing classes of matroids through the use of excluded or forbidden minors is an important and complicated aspect of matroid theory for which many interesting open problems remain [Oxl92]. Besides providing a uniform framework to study families of (Δ-)matroids, the introduction of sets of forbidden minors possesses other advantages. For example it allows one to compare easily between classes of (Δ-)matroids, especially when the complete set of minimal excluded minors is finite and known. We adopt this technique here to show that the four classes of Δ-matroids we are concerned with, compact, co-independent, binary, and local, are pairwise incomparable.

We begin with a brief introduction of forbidden minors. Let $S = (E, \mathcal{F})$ be a Δ-matroid. For $u \in E$, we let $S_{\overline{u}} = (E \setminus \{u\}, \{F \subseteq E \setminus \{u\} : F \in \mathcal{F}\})$ and $S_u = (E \setminus \{u\}, \{F \subseteq E \setminus \{u\} : F \cup \{u\} \in \mathcal{F}\})$. We call $S_{\overline{u}}$ and S_u the *elementary minors* of S at u. A *minor* is obtained by taking successive elementary minors. Let $\mathcal{S}$ be a set of Δ-matroids. We say that a Δ-matroid O is a *forbidden minor* of $\mathcal{S}$ if every Δ-matroid that has a minor isomorphic to O does not belong to $\mathcal{S}$, if, in addition, every minor of O is in $\mathcal{S}$ then O is a *minimal forbidden minor* of $\mathcal{S}$. The class of all minimal forbidden minors, up to isomorphism, of $\mathcal{S}$ defines an *obstruction set* of $\mathcal{S}$. Some of the families of Δ-matroids introduced in these pages are known to have simple obstruction sets.

The obstruction set for the class $\mathcal{S}B$ of binary Δ-matroids is $\{B_1, \ldots, B_5\}$ where $B_1 = (\{1,2,3\}, \{\emptyset, 12, 23, 31, 123\}), B_2 = (\{1,2,3\}, \{\emptyset, 1, 2, 3, 12, 23, 31\}), B_3 = (\{1,2,3\}, \{\emptyset, 2, 3, 13, 12, 123\}),$ $B_4 = (\{1,2,3,4\}, \{\emptyset, 12, 13, 14, 23, 24, 34\}), B_5 = (\{1,2,3,4\}, \{\emptyset, 12, 23, 34, 41, 1234\})$ [BD91].

It is easy to verify that the minimal obstruction set for the class $\mathcal{S}C$ of co-independent Δ-matroids is $\{C_1\}$, where $C_1 = (\{1\}, \emptyset)$

The obstruction set for the class $\mathcal{S}L$ of local Δ-matroids is $\{L_1, \ldots, L_5\}$ where $L_1 = (\{1,2,3\}, \{\emptyset, 2, 3, 12, 13\})$, $L_2 = (\{1,2,3\}, \{\emptyset, 3, 12, 13, 23\})$, $L_3 = (\{1,2,3\}, \{\emptyset, 2, 3, 12, 13, 23\})$, $L_4 = (\{1,2,3,4\}, \{\emptyset, 12, 13, 24, 34\})$, and $L_5 = (\{1,2,3,4\}, \{\emptyset, 12, 13, 23, 24, 34\})$ (see Appendix E).

We do not know the obstruction set for the class $\mathcal{S}O$ of compact Δ-matroids. We do, however, know some forbidden minors for compact Δ-matroids. In particular one can verify that the following are forbidden minors for $\mathcal{S}O$: $O_1 = (\{1,2,3,4\}, \mathcal{P}(\{1,2,3,4\}) \setminus \{\emptyset, 34, 1234\})$ and $O_2 = (\{1,2,3,4,5\}, \{13, 14, 15, 23, 24, 25\})$.

The following table is indexed by the four classes of Δ-matroids: $\mathcal{S}O, \mathcal{S}C, \mathcal{S}B, \mathcal{S}L$. The table is filled with Δ-matroids in such a way that if if S is the non-diagonal entry of the table corresponding to the row indexed by $\mathcal{S}_1$ and to the column indexed by $\mathcal{S}_2$ then $S \in \mathcal{S}_1$ (easily proved from the definition of the class) and $S \notin \mathcal{S}_2$ (since S is a forbidden minor for $\mathcal{S}_2$).

	$\mathcal{S}O$	$\mathcal{S}C$	$\mathcal{S}B$	$\mathcal{S}L$
$\mathcal{S}O$		C_1	B_2	L_3
$\mathcal{S}C$	O_1		B_2	L_3
$\mathcal{S}B$	O_2	C_1		L_4
$\mathcal{S}L$	O_1	C_1	B_5	

In consequence, every pair of families of Δ-matroids is pairwise incomparable.

Furthermore, if $\mathcal{S}$ is a collection of Δ-matroids that possesses an obstruction set of finite cardinality then we can decide whether a Δ-matroid S belongs to $\mathcal{S}$ by computing every minor of S and verifying whether it is isomorphic to a member of the obstruction set. In consequence we have:

Claim 4.1. Let $\mathcal{S}$ be a class of Δ-matroids that possess an obstruction set of finite cardinality. Then membership in $\mathcal{S}$ is decidable.

Corollary 4.2. *The problem of determining whether a given basis is local (binary, co-independent) is decidable.*

References

[BC95] André Bouchet and William H. Cunningham. Delta-matroids, jump systems, and bisubmodular polyhedra. *SIAM J. Discrete Math.*, 8(1):17–32, 1995.

[BD91] A. Bouchet and A. Duchamp. Representability of $\triangle$-matroids over gf(2). *Linear Algebra Appl.*, 146:67–78, 1991.

[Ber57] Claude Berge. Two theorems in graph theory. *Proc. Nat. Acad. Sci. U.S.A.*, 43:842–844, 1957.

[BJ00] André Bouchet and Bill Jackson. Parity systems and the delta-matroid intersection problem. *Electron. J. Combin.*, 7(1):Research Paper 14, 22 pp. (electronic), 2000.

[Bou88] A. Bouchet. Representability of $\triangle$-matroids. In *Combinatorics (Eger, 1987)*, pages 167–182. North-Holland, Amsterdam, 1988.

[Bou89] André Bouchet. Matchings and $\triangle$-matroids. *Discrete Appl. Math.*, 24(1-3):55–62, 1989. First Montreal Conference on Combinatorics and Computer Science, 1987.

[Bou95] André Bouchet. Coverings and delta-coverings. In *Integer programming and combinatorial optimization (Copenhagen, 1995)*, pages 228–243. Springer, Berlin, 1995.

[CKS01] Nadia Creignou, Sanjeev Khanna, and Madhu Sudan. *Complexity classifications of Boolean constraint satisfaction problems*. SIAM Monographs on Discrete Mathematics and Applications. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.

[Edm65] Jack Edmonds. Paths, trees, and flowers. *Canad. J. Math.*, 17:449–467, 1965.

[Fed01] Tomás Feder. Fanout limitations on constraint systems. *Theoret. Comput. Sci.*, 255(1-2):281–293, 2001.

[GIM97] J. Geelen, S. Iwata, and K. Murota. The linear delta-matroid parity problem. Technical Report RIMS Preprint 1149, Kyoto University, 1997.

[GJ79] Michael R. Garey and David S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.

[Ist97] Gabriel I. Istrate. Looking for a version of schaeer's dichotomy theorem when each variable occurs at most twice. Technical Report TR652, The University of Rochester, March 1997.

[KST93] Jan Kratochvíl, Petr Savický, and Zsolt Tuza. One more occurrence of variables makes satisfiability jump from trivial to NP-complete. *SIAM J. Comput.*, 22(1):203–210, 1993.

[Lad75] Richard E. Ladner. On the structure of polynomial time reducibility. *J. Assoc. Comput. Mach.*, 22:155–171, 1975.

[Lov81] L. Lovász. The matroid matching problem. In *Algebraic methods in graph theory, Vol. I, II (Szeged, 1978)*, pages 495–517. North-Holland, Amsterdam, 1981.

[Lov97] László Lovász. The membership problem in jump systems. *J. Combin. Theory Ser. B*, 70(1):45–66, 1997.

[LP86] L. Lovász and M. D. Plummer. *Matching Theory*. North-Holland Publishing Co., Amsterdam, 1986. Annals of Discrete Mathematics, 29.

[Oxl92] James G. Oxley. *Matroid theory*. The Clarendon Press Oxford University Press, New York, 1992.

[Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Conference Record of the Tenth Annual ACM Symposium on Theory of Computing (San Diego, Calif., 1978)*, pages 216–226. ACM, New York, 1978.

[Wen99] W. Wenzel. A unified treatment of the geometric algebra of matroids and even delta-matroids. *ADVANCES IN APPLIED MATHEMATICS*, 22(4):413–453, MAY 1999.

Appendix A. Proof of Theorem 2.3

Proof. We only have to show that $SAT_C(\Gamma)$ is no harder than $SAT_C(k,\Gamma)$ when Γ is non-Schaefer and $k \geq 3$. We use intensively the concept of *representable relation* [CKS01] which has been central in the study of generalized satisfiability problems. However, in order to prove Theorem 2.3 we need a stronger notion of representability that takes into account the maximum number of occurrences per variable. The basic idea is the following, let ϕ be any $CNF(k,\Gamma)$-formula and select a set of variables $v_1,\ldots,v_m$ occurring once in ϕ. The set of assignments to the variables $v_1,\ldots,v_m$ such that there exists a valid extension satisfying ϕ defines a relation R_ϕ. We call every relation obtained this way *k-representable*. Notice that if we add the relation R_ϕ to Γ we do not increase the complexity of $SAT_C(k,\Gamma)$ since every occurrence of a predicate R'_ϕ interpreted as the relation, R_ϕ, can be replaced with ϕ (duplicating the auxiliary variables if necessary).

We will abuse notation by occasionally denoting a relation by some equivalent logical expression. For example we will denote relation $R := \{(0,1),(1,0),(1,1)\}$ by $x \vee y$. We will also denote by $x \vee y$ the clause $R'(x,y)$.

The proof of the theorem is built upon some basic results. In the following $k \geq 3$.

• If there is a relation $R \in \Gamma$ which is not Horn, then either $x \vee y$ or $x \neq y$ can be k-represented. By duality, if there is a relation $R \in \Gamma$ which is not dual-Horn, then either $\overline{x} \vee \overline{y}$ or $x \neq y$ can be k-represented.

Proof. Since R is not Horn there exists a minimal subset, $I = \{i_1,\ldots,i_j\}$, $1 \leq i_1 < \cdots < i_j \leq r$ where r is the arity of R, such that the *projection* R_I defined as $\{(\mathbf{t}[i_1],\ldots,\mathbf{t}[i_j]) | \mathbf{t} \in R\}$, is not Horn. Clearly R_I can be k-represented. Since R_I is not horn, there exists two tuples $\mathbf{t_1},\mathbf{t_2} \in R_I$ such that $\mathbf{t} = \mathbf{t_1} \wedge \mathbf{t_2}$ is not in R_I. First, notice that there exists some index $1 \leq l_1 \leq j$ such that $\mathbf{t_1}[l_1] = 1$ and $\mathbf{t_2}[l_1] = 0$ since otherwise $\mathbf{t_1} \wedge \mathbf{t_2} = \mathbf{t_1}$ which is in R_I. Similarly, we can conclude that there exists some index $1 \leq l_2 \leq j$ such that $\mathbf{t_1}[l_2] = 0$ and $\mathbf{t_2}[l_2] = 1$. We then construct the formula $\phi(x,y) = R'_I(v_1,\ldots,v_j)$ where $v_{l_1} = x$, $v_{l_2} = y$, $v_n = 1$ if $\mathbf{t_1}[n] \wedge \mathbf{t_2}[n] = \mathbf{1}$ and $v_n = \mathbf{0}$ otherwise, $n \in \{1,\ldots,r\}/\{l_1,l_2\}$.

Let us study the relation R_ϕ k-represented by formula ϕ. First, since $\mathbf{t} \notin R_I$ we have that $(0,0) \notin R_\phi$. Next, we will see that $(0,1) \in R_\phi$ since otherwise $R^* = (R_I)_{\{1,\ldots,r\}/\{l_1\}}$ is not Horn, contradicting the minimality of I. To show that relation R^* is not Horn, we observe that the projections of $\mathbf{t_1}$ and $\mathbf{t_2}$ to $\{1,\ldots,r\}/\{l_1\}$ belong to R^* whereas its intersection does not belong to it. Similarly, we show that $(1,0) \in R_\phi$. Thus, two possibilities remain: if $(1,1) \in R_\phi$ then $\phi(x,y)$ k-represents $x \vee y$ whereas if $(1,1) \notin R_\phi$ then $\phi(x,y)$ k-represents $x \neq y$. □

•If Γ is neither Horn nor dual-Horn, then $x \neq y$ or $\overline{x} \vee y$ can be k-represented.

Proof. If Γ is neither Horn nor dual-Horn, then we have that either $x \neq y$ can be k-represented or both $x \vee y$ or $\overline{x} \vee \overline{y}$ can be k-represented. In the first case we are done. For the second we construct the formula $\phi = (\overline{x} \vee \overline{z}) \wedge (z \vee y)$, which over variables x, y k-represents relation $\overline{x} \vee y$. □

To conclude the proof of the Theorem, let Γ be a non-Schaefer basis. Then either $x \neq y$ or $\overline{x} \vee y$ can be k-represented. Thus, for every $CNF(\Gamma)$-formula, ϕ, it is possible to construct a $CNF(3,\Gamma)$-formula, φ, replacing all the occurrences of a variable x with different copies of the variables $x_1,\ldots,x_n$ and using relation $x \neq y$ or $\overline{x} \vee y$ to enforce all the variables, $x_1,\ldots,x_n$, to take the same value. If $x \neq y$ can be k-represented then φ would contain $(x_1 \neq y_1) \wedge (y_1 \neq x_2) \wedge \cdots \wedge (y_{n-1} \neq x_n)$ where $y_1,\cdots,y_{n-1}$ are new variables, whereas if $\overline{x} \vee y$ can be k-represented then φ would contain $(\overline{x_1} \vee x_2) \wedge \cdots \wedge (\overline{x_{n-1}} \vee x_n) \wedge (\overline{x_n} \vee x_1)$. Notice that every variable in φ occurs at most 3 times. In both cases φ is satisfiable if and only if so is ϕ. Consequently, $SAT_C(k,\Gamma)$ is NP-complete. □

Appendix B. Proof of Theorems 3.4 and 3.15

Proof. Let ϕ be an instance of $SAT_C(2,\Gamma)$, that is, ϕ is a $CNF(\Gamma)$ formula. As in [Fed01] we can assume that every variable of ϕ occurs exactly twice in ϕ, otherwise we could construct a new $CNF(2,\Gamma)$ formula ψ in which every variable occurs exactly twice and such that ψ is satisfiable iff so is ϕ. The new formula ψ is defined to be the conjunction $\phi \wedge \phi'$ where ϕ' is a formula obtained from ϕ by replacing every variable occurring twice in ϕ by a new auxiliary variable.

We will construct an instance $((E,\mathcal{F}),\mathcal{L})$ of the $\mathcal{S}$-parity problem such that there exists some $F \in \mathcal{F}$ such that $\mathcal{L}_F = \mathcal{L}$ iff ϕ has a solution. Let φ be a $CNF(1,\Gamma)$ formula obtained from ϕ by replacing the two occurrences of a variable by two new variables, let E be the set of variables of φ and let $\mathcal{L}$ be the partition of E that pairs up every two variables in φ corresponding to the same variable in ϕ. The set $\mathcal{F}$ is defined to contain exactly those subsets of E such that their indicator function is a satisfying assignment of φ, that is, $\mathcal{F}$ is the set containing every $F \subseteq E$ such that $f : E \to \{0,1\}$ given by $f(v) = 1$ if $v \in F$ and 0 otherwise, is a satisfying assignment of φ. It is easy to observe that every feasible set F such that $\mathcal{L}_F = \mathcal{L}$ correspond to a satisfying assignment of φ that sets every pair of variables in $\mathcal{L}$ to the same value, and hence to a satisfying assignment of ϕ.

Now we will show that $(E, \mathcal{F})$ belongs to $\mathcal{S}$. Let $C = R'(x_1, \ldots, x_m)$ be a clause of φ and let $(E_C, \mathcal{F}_C)$ be the Δ-matroid with $E_C = \{x_1, \ldots, x_k\}$ and such that $\mathcal{F}_C$ contains exactly those subsets of E_C such that their indicator vector satisfies $R'(x_1, \ldots, x_k)$. Clearly, $(E_C, \mathcal{F}_C)$ belongs to $\mathcal{S}$, since it is isomorphic to R. Furthermore, if $\varphi = C_1 \wedge \cdots \wedge C_m$ then $(E, \mathcal{F})$ is the direct sum of all $(E_{C_i}, \mathcal{F}_{C_i}), 1 \le i \le m$. This completes the proof of Theorem 3.4.

In order to prove Theorem 3.15 we have to show additionally that a matrix representation of $(E, \mathcal{F})$ can be computed in polynomial time. Since the basis Γ is finite we can assume that a matrix representation for every relation in Γ (and hence for every $(E_C, \mathcal{F}_C)$) is known. So we only have to show how, given a family Δ-matroids specified by means of their matrix representation we can compute the matrix representation of its direct sum. This part is better viewed in terms of relations than of set systems. Thus, let $R_1, \ldots, R_m$ be binary Δ-matroids. For each $R_i, 1 \le i \le n$ we denote by $(A_i, \mathbf{x}_i)$ its matrix representation. Recall that the direct sum of $R_1, \ldots, R_m$ corresponds to the Cartesian product $R = R_1 \times \cdots \times R_m$.

Define $\mathbf{x} := (\mathbf{x}_1, \ldots, \mathbf{x}_m)$ and M to be the block diagonal matrix with blocks $M_1, \ldots, M_m$. Then, it suffices to prove the claim that $(M, \mathbf{x})$ is a matrix representation of R.

Note that M is a symmetric matrix over $GF(2)$ since each of the M_i are symmetric matrices over $GF(2)$, and block diagonal matrices consisting of symmetric blocks are also symmetric.

We want to prove that for every $\mathbf{t}$, $\mathbf{t} \in A \oplus \mathbf{x}$ if and only if rank $M[\mathbf{t}] = \|\mathbf{t}\|_1$. Let $\mathbf{t} = (\mathbf{t}_1, \ldots, \mathbf{t}_m)$. By construction, $\mathbf{t}_i \in A_i \oplus \mathbf{x}_i$ if and only if rank $M_i[\mathbf{t}_i] = \|\mathbf{t}_i\|_1$, for $i = 1, \ldots, m$. So that $(\mathbf{t}_1, \ldots, \mathbf{t}_m) \in A \oplus \mathbf{x}$ if and only if rank $M[\mathbf{t}] = \sum_{1 \le i \le m}$ rank $M_i[\mathbf{t}_i] = \sum_{1 \le i \le m} \|\mathbf{t}_i\|_1 = \|\mathbf{t}\|_1$.

This completes the proof of Theorem 3.15 ☐

Appendix C. Augmenting Paths

Theorem. *Let $((E, \mathcal{F}), \mathcal{L})$ be an instance of the Δ-matroid parity problem and $F \in \mathcal{F}$ such that $|\mathcal{L}_F|$ is not maximum. Then there exists an F-augmenting $\mathcal{L}_F$-alternating path. Furthermore, we can compute it in time polynomial in $|E|$ given a $G \in \mathcal{F}$ with $|\mathcal{L}_G| > |\mathcal{L}_F|$.*

Proof. We will provide an algorithm that finds an F-augmenting $\mathcal{L}_F$-alternating path. The algorithm has two procedures. The first procedure *Find-Closer-Not-Smaller (FCNS)* receives as input two feasible sets $F, G \in \mathcal{F}$ and computes, in time polynomial in $|E|$, some $G^* \in \mathcal{F}$ such that $\mathcal{L}_F \subseteq \mathcal{L}_{G^*}$ and $|\mathcal{L}_G| \le |\mathcal{L}_{G^*}|$.
Procedure Find-Closer-Not-Smaller (FCNS) (F, G)

- **Step 1:** If $\mathcal{L}_F \subseteq \mathcal{L}_G$ then return G.
- **Step 2:** Find some $u \in F \Delta G$ such that $\{u, \overline{u}\} \in \mathcal{L}_F \setminus \mathcal{L}_G$.
- **Step 3:** Find some $v \in F \Delta G$ such that $G \Delta \{u, v\} \in \mathcal{F}$.
- **Step 4:** G^*:=Find-Closer-Not-Smaller $(F, G\Delta\{u, v\})$. Return G^*.

Let us analyze the procedure. We will show by induction on $s = |F \Delta G|$ that: (A) the procedure FCNS finds some G^* such that $\mathcal{L}_F \subseteq \mathcal{L}_{G^*}$ and $|\mathcal{L}_G| \le |\mathcal{L}_{G^*}|$. Statement (A) is clearly true for $s = 0$, since in this case procedure FCNS stops at step 1 returning G. Assume now that statement (A) is true for every $F, G \in \mathcal{F}$ such that $|F \Delta G| \le s, s \ge 0$ and let $F, G \in \mathcal{F}$ such that $|F \Delta G| = s + 1$. If $\mathcal{L}_F \subseteq \mathcal{L}_G$ then procedure FCNS returns G which is a valid answer. Otherwise, the procedure calls recursively procedure FCNS$(F, G\Delta\{u, v\})$ and returns the output of the call. Since $|F \Delta (G \Delta \{u, v\})| \le s$, by induction hypothesis the call returns some G^* such that $\mathcal{L}_F \subseteq \mathcal{L}_{G^*}$ and $|\mathcal{L}_{G\Delta\{u,v\}}| \le |\mathcal{L}_{G^*}|$. Thus, in order to show that G^* satisfies (A) we only have to show that $|\mathcal{L}_G| \le |\mathcal{L}_{G\Delta\{u,v\}}|$. In order to see this we do a case analysis: Let $G' = G\Delta\{u, v\}$. If $v = u$, then $\mathcal{L}_{G'} = \mathcal{L}_G \cup \{\{u, \overline{u}\}\}$ and hence $|\mathcal{L}_{G'}| = |\mathcal{L}_G| + 1$. Otherwise $\mathcal{L}_{G'} = (\mathcal{L}_G \cup \{\{u, \overline{u}\}\})/\{\{v, \overline{v}\}\}$, or $\mathcal{L}_{G'} = \mathcal{L}_G \cup \{\{u, \overline{u}\}, \{v, \overline{v}\}\}$ depending on whether $\{v, \overline{v}\} \in \mathcal{L}_G$ or not. In both cases $|\mathcal{L}_{G'}| \ge |\mathcal{L}_G|$.

That shows the correctness of the procedure. In order to bound the complexity we just have to notice that since at every recursive call of the procedure $F \Delta G$ decreases. Thus we have at most $|E|$ recursive calls.

The second procedure, computes given $F, G \in \mathcal{F}$, such that $\mathcal{L}_F \subsetneq \mathcal{L}_G$ and $u \in F\Delta G$ such that $\{u, \overline{u}\} \in \mathcal{L}_G \setminus \mathcal{L}_F$ an F-augmenting $\mathcal{L}_F$-alternating path $u_1, \ldots, u_n$ such that $\{u_1, \ldots, u_n\} \subseteq F\Delta G$ and $u_1 = u$.
procedure Find-Augmenting-Path (FAP) (u, F, G)

Step 1: Find v such that $F\Delta\{u, v\} \in \mathcal{F}$.
Step 2: If $v = u$ then return u.
Step 3: If $\{v, \overline{v}\} \notin \mathcal{L}_F$ then return u, v.
Step 4: $v_1, \ldots, v_n$:=Find-Augmenting-Path $(\overline{v}, F\Delta\{u, v\}, G)$.
Step 5: Return $u, v, v_1, \ldots, v_n$

In order to show the correctness of the procedure we will start by observing that if procedure FAP stops at step 2 or step 3, then it provides a valid answer.

We will show by induction on $s = |F\Delta G|$, that for every $F, G \in \mathcal{F}$, such that $\mathcal{L}_F \subsetneq \mathcal{L}_G$, and every $u \in F\Delta G$ such that $\{u, \overline{u}\} \in \mathcal{L}_G \setminus \mathcal{L}_F$: (B) procedure FAP computes an F-augmenting $\mathcal{L}_F$-alternating path $u_1, \ldots, u_n$ such that $\{u_1, \ldots, u_n\} \subseteq F\Delta G$ and $u_1 = u$.

Statement (B) is true for $s = 1$, since in this case $v = u$ and consequently FAP stops at step 2 returning v which is a valid answer. Now assume that the statement (B) holds for every $F, G \in \mathcal{F}$ such that $|F\Delta G| \leq s$, let $F, G \in \mathcal{F}$ such that $|F\Delta G| = s + 1$, $\mathcal{L}_F \subsetneq \mathcal{L}_G$, and let $u \in F\Delta G$ such that $\{u, \overline{u}\} \in \mathcal{L}_G \setminus \mathcal{L}_F$. If procedure FAP stops at steps 2 or 3 then it provides a valid answer. Otherwise, at step 4 it calls recursively FAP with parameters $\overline{v}$, $F' = F\Delta\{u, v\}$, and G. Notice that $|F'\Delta G| \leq s$ so, by induction hypothesis, we can assume that the output $v_1, \ldots, v_n$ is an F'-augmenting $\mathcal{L}_{F'}$-alternating path such that $v_1 = \overline{v}$ and $\{v_1, \ldots, v_n\} \subseteq F'\Delta G$. We will show that the path $u, v, v_1, \ldots, v_n$ satisfies the required conditions. Since $u, v \in F\Delta G$ and $\{v_1, \ldots, v_n\} \in F'\Delta G \subseteq F\Delta G$, we have that $\{u, v, v_1, \ldots, v_n\} \subseteq F\Delta G$. Since $F\Delta\{u, v\} \in \mathcal{F}$ and $v_1, \ldots, v_n$ is $F\Delta\{u, v\}$-augmenting we have that $u, v, v_1, \ldots, v_n$ is F-augmenting. It only remains to show that $u, v, v_1, \ldots, v_n$ satisfies the tree properties necessary to be an $\mathcal{L}_F$-alternating path. First notice that $\mathcal{L}_{F'} = (\mathcal{L}_F \cup \{u, \overline{u}\}) \setminus \{v, \overline{v}\}$ and furthermore for every $2 \leq l \leq n$, $v_l \notin \{u, \overline{u}, v, \overline{v}\}$. In consequence $\{v_l, \overline{v_l}\} \in \mathcal{L}_F$ iff $\{v_l, \overline{v_l}\} \in \mathcal{L}_{F'}$. With this observation it is an easy task to show that $u, v, v_1, \ldots, v_n$ satisfies conditions (1-3) of an $\mathcal{L}_F$-alternating path. Condition 1: Since algorithm reaches step 4, we know that $\{v, v_1\} \in \mathcal{L}_F$. Furthermore for every $1 \leq 2j < n$, $\{v_{2j}, v_{2j+1}\}$ is in $\mathcal{L}_{F'}$, by induction hypothesis, and hence is is also in $\mathcal{L}_F$. Condition 2 is trivially satisfied since by hypothesis $\{u, \overline{u}\} \notin \mathcal{L}_F$. Condition 3: by hypothesis $\overline{u} \notin F\Delta G$, thus since $v_n \in F'\Delta G \subseteq F\Delta G$, we can conclude that $\overline{u} \neq v_n$. Furthermore, if n is even by induction hypothesis $\{v_n, \overline{v_n}\}$ is not in $\mathcal{L}_{F'}$ and hence is not in $\mathcal{L}_F$.

In order to bound the time complexity we observe that at every recursive call of the procedure, $F\Delta G$ is decreased. Thus, the number of calls is at most $|E|$.

With the help of procedures FCNS and FAP it is an easy task to design an algorithm that computes, given some $F, G \in \mathcal{F}$, such that $|\mathcal{L}_F| \subseteq |\mathcal{L}_G|$, an F-augmenting $\mathcal{L}_F$-alternating path.
Algorithm (F, G)

Step 1: G^*:=Find-Closest-Not-Smaller (F, G)
Step 2: Find some $u \in F\Delta G^*$ such that $\{u, \overline{u}\} \in \mathcal{L}_{G^*} \setminus \mathcal{L}_F$.
Step 3: $u_1, \ldots, u_n$:=Find-Augmenting-Path (u, F, G).
Step 4: Return $u_1, \ldots, u_n$.

□

Corollary. *For every collection $\mathcal{S}$ of Δ-matroids, the $\mathcal{S}$-augmenting path problem is polynomially equivalent to the $\mathcal{S}$-parity problem.*

Proof. In order to show that if the $\mathcal{S}$-parity problem is polynomially solvable then so is the $\mathcal{S}$-augmenting problem we have only to appeal to Theorem 3.6.

For the converse we will start by showing an intermediate result that will be very useful.

Let $((E, \mathcal{F}), \mathcal{L})$ be an instance of the Δ-matroid, let $F \in \mathcal{F}$ be a feasible set, let $u_1, \ldots, u_n$ be a F-augmenting $\mathcal{L}_F$-alternating path, and let F' be $F\Delta\{u_1, \ldots, u_n\}$. Then we have that $F' \in \mathcal{F}$, since $u_1, \ldots, u_n$ is an F-augmenting path. Furthermore, since $u_1, \ldots, u_n$ is an $\mathcal{L}_F$-alternating path we have that $\mathcal{L}'_F = \mathcal{L}_F \cup \{\{u_1, \overline{u_1}\}\}$ if n is odd and $\mathcal{L}'_F = \mathcal{L}_F \cup \{\{u_1, \overline{u_1}\}, \{u_n, \overline{u_n}\}\}$ if n even. Consequently in both cases $|\mathcal{L}_{F'}| > |\mathcal{L}_F|$.

Thus, it is easy to design an algorithm that solves the $\mathcal{S}$-parity problem in polynomial time provided that the $\mathcal{S}$-augmenting problem is polynomially solvable:

Start with an arbitrary $F \in \mathcal{F}$. If there not exists any F-augmenting $\mathcal{L}_F$-alternating path, then $\mathcal{L}_F$ is maximum and hence we are done. Otherwise, let $u_1, \ldots, u_n$ be such a path. Then $F' = F\Delta\{u_1, \ldots, u_n\}$ also belongs to $\mathcal{F}$ and furthermore $|\mathcal{L}_{F'}| > |\mathcal{L}_F|$. With at most $|E|/2$ iterations of this step we will find an optimum feasible set. □

Appendix D. Proof of Theorem 3.10

The proof of Theorem 3.10 uses several intermediate results. Basically we show that over local Δ-matroids augmenting paths and local paths are equivalent in the sense that one can be easily obtained from the other. We then show that the problem of finding a local path can be solved in polynomial time.

We start by showing that when we are dealing with local Δ-matroids, every minimal augmenting path is also local and that every minimal local path is also augmenting. We actually prove that this is true even for 2-local Δ-matroids.

We need a last definition. Let $(E, \mathcal{F})$ be a Δ-matroid and let $F \in \mathcal{F}$. A path is F-augmenting-minimal if it is F-augmenting and it does not have any proper F-augmenting subpath.

Lemma D.1. *Let $(E,\mathcal{F})$ be a 2-local Δ-matroid and let F be any element in $\mathcal{F}$. Every F-augmenting-minimal path is also F-local.*

Proof. We shall show by induction on m that every F-augmenting-minimal path of length at most $2m$ is also F-local. The statement clearly holds for $m = 1$. For the induction step assume that the statement holds for m and let $u_1, \ldots, u_{2m+1}$ be a F-augmenting-minimal path of length $2m+1$. Notice that if $u_1, \ldots, u_{2m+1}$ is a F-augmenting-minimal path then $u_1, \ldots, u_{2m}$ is a F-augmenting-minimal path of length $2m$ and consequently, $F\Delta\{u_{2j-1}, u_{2j}\} \in \mathcal{F}$, for all $1 \leq 2j \leq 2m$ by induction hypothesis.

Thus, it suffices to show that $F\Delta\{u_{2m+1}\} \in \mathcal{F}$. In order to do this we will show by (backwards) induction that for $0 \leq l \leq m$, $F\Delta\{u_1, \ldots, u_{2l}, u_{2m+1}\} \in \mathcal{F}$. The statement is clearly true for $l = m$. Now we will show that the statement holds for $l - 1$ whenever it holds for l. First, notice that $u_1, \ldots, u_{2l}$ is an F-augmenting path and thus, $F\Delta\{u_1, \ldots, u_{2l-2}\}$ and $F\Delta\{u_1, \ldots, u_{2l}\}$ are in $\mathcal{F}$. Thus, $u_{2l}, u_{2l-1}, u_{2m+1}$ is a $F\Delta\{u_1, \ldots, u_{2l}\}$-local path. If it is local-minimal then it is also an $F\Delta\{u_1, \ldots, u_{2l}\}$-augmenting path and hence $F\Delta\{u_1, \ldots, u_{2l-2}, u_{2m+1}\} \in \mathcal{F}$ and we are done. Otherwise u_{2l} is an $F\Delta\{u_1, \ldots, u_{2l}\}$-augmenting path and hence $u_1, \ldots, u_{2l-1}$ is a F-augmenting subpath of $u_1, \ldots, u_{2m+1}$ in contradiction with the minimality assumption. A similar argument applies for F-augmenting-minimal paths $u_1, \ldots, u_{2m+2}$ of length $2m + 2$.

We show that $F\Delta\{u_{2m+1}, u_{2m+2}\} \in \mathcal{F}$ by showing that for $1 \leq l \leq m$, $F\Delta\{u_1, \ldots, u_{2l}, u_{2m+1}, u_{2m+2}\} \in \mathcal{F}$ implies that $F\Delta\{u_1, \ldots, u_{2l-2}, u_{2m+1}, u_{2m+2}\} \in \mathcal{F}$. Since $u_1, \ldots, u_{2l}$ is an F-augmenting path, then we have that $F\Delta\{u_1, \ldots, u_{2l-2}\}$ and $F\Delta\{u_1, \ldots, u_{2l}\}$ are in $\mathcal{F}$. In consequence we have that $u_{2l}, u_{2l-1}, u_{2m+1}, u_{2m+2}$ is a $F\Delta\{u_1, \ldots, u_{2l}\}$-local path whose only possible augmenting-minimal subpaths are the following: (1) $u_{2l}, u_{2l-1}, u_{2m+1}, u_{2m+2}$, (2) u_{2l}, (3), $u_{2l}, u_{2l-1}, u_{2m+1}$, and (4) u_{2l}, u_{2m+2}.

In case (1) $F\Delta\{u_1, \ldots, u_{2l-2}, u_{2m+1}, u_{2m+2}\} \in \mathcal{F}$ and we are done. In case (2) $u_1, \ldots, u_{2l-1}$ is a F-augmenting subpath of $u_1, \ldots, u_{2m+2}$, reaching a contradiction. In case (3) $F\Delta\{u_1, \ldots, u_{2l-2}, u_{2m+1}\} \in \mathcal{F}$. We will show that $F\Delta\{u_1, \ldots, u_{2m}, u_{2m+1}\} \in \mathcal{F}$. In consequence, $u_1, \ldots, u_{2m+1}$ is a F-augmenting subpath of $u_1, \ldots, u_{2m+2}$, a contradiction.

To do it we will show by induction that for $l-1 \leq j \leq m$, $F\Delta\{u_1, \ldots, u_{2j}, u_{2m+1}\}$. The statement is clearly true for $j = l - 1$. Assume now that is is true for j. Thus, $u_{2j+1}, u_{2j+2}, u_{2m+1}$ is a $F\Delta\{u_1, \ldots, u_{2j}\}$-local path. If it is is local-minimal, then $F\Delta\{u_1, \ldots, u_{2j+1}, u_{2j+2}, u_{2m+1}\} \in \mathcal{F}$ and we are done. Otherwise, u_{2j+1} is local-minimal and consequently $u_1, \ldots, u_{2j+1}$ is a F-augmenting subpath of $u_1, \ldots, u_{2m+1}$, a contradiction again. Finally, in case (4) $u_1, \ldots, u_{2l-1}, u_{2m+2}$ is a F-augmenting subpath of $u_1, \ldots, u_{2m+2}$, a contradiction. □

Lemma D.2. *Let $(E,\mathcal{F})$ be a Δ-matroid, given as a feasible oracle, let $\mathcal{L}$ be a partition of E into pairs, and let $F \in \mathcal{F}$. It is possible to decide whether there exists F-local $\mathcal{L}_F$-alternating path and compute a F-local-minimal one, if exists, in time polynomial in $|E|$.*

Proof. We will reduce it to the of computing a F-local $\mathcal{L}_F$-alternating path to the well studied problem of finding an augmenting path for a matching in a graph.

We start by recalling some basic concepts about matching and graphs. Let $G = (V, H)$ be a non-directed graph. A matching G is any subset M of H such that no two edges of M share a node. Every edge of M

is called matched whereas every edge in H/E is called free. Every node not incident with any edge of M is called exposed. An augmenting path for M is a path in G whose edges are alternatively matched and free and such that both its end nodes are exposed. Furthermore, Edmond's algorithm [Edm65] decides whether a matching has an augmenting path, and finds one, if one exists, in time polynomial in the size of H.

Let U be the set containing all the elements in E paired by F, that is $U = \{u : \{u, \overline{u}\} \in \mathcal{L}_F\}$, and let u and u' be two elements in E/U such that $\{u, u'\} \notin \mathcal{L}$. We will construct a graph G and a matching M such that M has an augmenting path if and only if there exists an F-local $\mathcal{L}_F$-alternating path $u_1, \ldots, u_{2m}$, $m \geq 1$ such that $u_1 = u$ and $u_{2m} = u'$.

We define G to be the graph (V, H) with $V = U \cup \{u, u'\}$ and set of edges H given by

$$\{\{x, y\} : x, y \in V, x \neq y, F\Delta\{x, y\} \in \mathcal{F}\} \cup \mathcal{L}_F.$$

Let M be the matching of G given by $M = \mathcal{L}_F$.

>From the way G and M are defined it is immediate to see that $u_1, \ldots, u_{2m}$, $m \geq 1$ iff an augmenting path for M iff and only iff it is an F-local $\mathcal{L}_F$-alternating path such that either $u_1 = u$ and $u_{2m} = u'$, or $u_1 = u'$ and $u_m = u'$. In the second case, we only need to reverse path obtaining again a F-local $\mathcal{L}_F$-alternating path.

Thus, in order to decide whether an F-local $\mathcal{L}_F$-alternating path of even length exists we only have to construct the previous graph for every possible pair u, u' and apply Edmond's algorithm to find an augmenting path.

We have also to take care of F-local paths $\mathcal{L}_F$-alternating paths of odd length. Let now u be any element in E/U. We will construct a graph G such and a matching M such that M has an augmenting path if and only if there exists an F-local path $\mathcal{L}_F$-alternating path $u_1, \ldots, u_{2m-1}, m \geq 1$ such that $u_1 = u$.

Let u' be any element not in E. We define G to be the graph (V, H) with $V = U \cup \{u, u'\}$ and set of edges H given by

$$\{\{x, y\} : x, y \in V, x \neq y, F\Delta\{x, y\} \in \mathcal{F}\} \cup \mathcal{L}_F \cup \{\{x, u'\} : x \in V, F\Delta\{x\} \in \mathcal{F}\}.$$

It is immediate to see that $u_1, \ldots, u_{2m-1}, u'$ is an augmenting path for M whenever $u = u_1, \ldots, u_{2m-1}$, $m \geq 1$ is a F-local $\mathcal{L}_F$-alternating path. For the converse, let $u_1, \ldots, u_{2m}$, $m \geq 1$ be a augmenting path for M. Thus clearly, since u and u' are the only exposed nodes of G, they must be the end nodes of the path. We can assume that $u_1 = u$ and $u_{2m} = u'$, since otherwise we could reverse the path obtaining again an augmenting path. Finally, it is immediate to see that $u_1, \ldots, u_{2m-1}$ is a F-local $\mathcal{L}_F$-alternating path.

Similarly to the previous case, we can decide whether there exists an F-local $\mathcal{L}_F$-alternating path of odd length by constructing the previous graph for every u not in E and applying Edmond's algorithm.

In order to complete the proof we have to see that we can compute a F-local-minimal path in polynomial time once a F-local $v_1, \ldots, v_n$ path is found. The procedure is simple, we only have to check whether there exists a cut or shortcut of $v_1, \ldots, v_n$ which is also F-local (we have a polynomial number of them). If every cut and shortcut is not F-local then $v_1, \ldots, v_n$ is local-minimal and we are done. Otherwise, we can find a F-local subpath of $v_1, \ldots, v_n$ from which we can iterate the process. In at most $n - 1$ iterations we must obtain a F-local-minimal path. □

Finally, to prove of Theorem 3.10 we only have to put together all the pieces. More precisely we have to show that given a local Δ-matroid $(E, \mathcal{F})$ a partition $\mathcal{L}$ of E into pairs and $F \in \mathcal{F}$, it is possible to find a F-augmenting $\mathcal{L}_F$-alternating path or determine that none exists in time polynomial in $|E|$ and then appeal to Theorem 3.7.

In order to do this we use the following observation of easy proof: Let $L \subseteq \mathcal{L}$ be any collection of pairs from $\mathcal{L}$. Every subpath of an L-alternating path it is also L-alternating.

The procedure to accomplish this goal is the following. By Lemma D.2 we can decide whether there exists a F-local $\mathcal{L}_F$-alternating path and find a minimal one, if exists, in polynomial time. If the algorithm provides us with a F-local-minimal $\mathcal{L}_F$-alternating path then we are done, since it is also F-augmenting. Otherwise, if we found out that such a path does not exist then we can conclude that there not exists any F-augmenting $\mathcal{L}_F$-alternating path, since if there exists a F-augmenting $\mathcal{L}_F$-alternating, then it has subpath which is also $\mathcal{L}_F$-alternating, F-augmenting-minimal, and, by Lemma D.1, F-local as well.

Appendix E. Some properties of local Δ-matroids

This appendix contains some basic properties satisfied by local Δ-matroids.

Lemma E.1. *The direct sum of local Δ-matroids is also local.*

Proof. Let $(E_1, \mathcal{F}_1)$ and $(E_2, \mathcal{F}_2)$ be local Δ-matroids. We will show that its direct sum $(E, \mathcal{F}) = (E_1 \cup E_2, \mathcal{F})$ is 2-local. Let F be an element of $\mathcal{F}$, thus, $F = F_1 \cup F_2$ where $F_1 \in \mathcal{F}_1$ and $F_2 \in \mathcal{F}_2$. Let u_1, u_2, u_3 be a F-local-minimal path. Assume wlog. that $u_1 \in E_1$. Notice that then u_2 must be also in E_1, since otherwise $F_1 \Delta \{u_1\} \in \mathcal{F}_1$, and hence u_1 would be F-local path, in contradiction with the F-local-minimality of u_1, u_2, u_3. Now we have two possibilities, if $u_3 \in E_1$, then $F_1 \Delta \{u_1, u_2, u_3\} \in \mathcal{F}_1$ and hence $F \Delta \{u_1, u_2, u_3\} \in \mathcal{F}$. Otherwise $F_2 \Delta \{u_3\} \in \mathcal{F}_2$, and consequently $F \Delta \{u_1, u_2, u_3\} = F_1 \Delta \{u_1, u_2\} \cup F_2 \Delta \{u_3\} \in \mathcal{F}$. Now, let u_1, u_2, u_3, u_4 be a F-local-minimal path. As before, we can assume wlog. that u_1 and u_2 are in E_1. Now we have two possibilities, if $u_3 \in E_1$, then u_4 must be also in E_1, since otherwise $F_1 \Delta \{u_3\} \in \mathcal{F}_1$, hence $F \Delta \{u_3\} \in \mathcal{F}$, and consequently u_1, u_2, u_3 would be an F-local subpath of u_1, u_2, u_3, u_4, a contradiction. Since u_3 and u_4 are also in E_1, we have that u_1, u_2, u_3, u_4 is a F_1-local path in $(E_1, \mathcal{F}_1)$ which should be also F_1-local-minimal. Thus, $F_1 \Delta \{u_1, u_2, u_3, u_4\} \in \mathcal{F}_1$ and consequently $F \Delta \{u_1, u_2, u_3, u_4\} \in \mathcal{F}$. In the case $u_3 \in E_2$, we can infer using a similar reasoning that u_4 must be also in E_2, and consequently $F_2 \Delta \{u_3, u_4\} \in \mathcal{F}_2$. Since $F_1 \Delta \{u_1, u_2\} \in \mathcal{F}_1$ we have that $F_1 \Delta \{u_1, u_2\} \cup F_2 \Delta \{u_3, u_4\} = F \Delta \{u_1, u_2, u_3, u_4\}$ must be in $\mathcal{F}$. □

Lemma E.2. *A Δ-matroid is local iff it is 2-local.*

Proof. We will show that if a Δ-matroid $(E, \mathcal{F})$ is M-local, $M \geq 2$, then it is also $(M+1)$-local. Let F be in $\mathcal{F}$ and let $u_1, \ldots, u_{2m+1}$ be a minimal F-local path of length $2m+1$. Since $u_1, \ldots, u_{2m}$ is also a F-local-minimal path we only have to show that $F \Delta \{u_1, \ldots, u_{2m}, u_{2m+1}\} \in \mathcal{F}$. In order to do this we will show by induction on l that for every $0 \leq l \leq m$, $F \Delta \{u_1, \ldots, u_{2l}, u_{2m+1}\} \in \mathcal{F}$. The statement is clearly true for $l = 0$. Assume now that it is true for $l < m$. Notice that $F \Delta \{u_1, \ldots, u_{2l}\}$, $F \Delta \{u_1, \ldots, u_{2l}, u_{2l+1}, u_{2l+2}\}$, and $F \Delta \{u_1, \ldots, u_{2l}, u_{2m+1}\}$ are in $\mathcal{F}$. Thus, $u_{2l+1}, u_{2l+2}, u_{2m+1}$ is a $F \Delta \{u_1, \ldots, u_{2l}\}$-local path. We have two possibilities: (1) $u_{2l+1}, u_{2l+2}, u_{2m+1}$ is local-minimal; consequently $F \Delta \{u_1, \ldots, u_{2l+2}, u_{2m+1}\} \in \mathcal{F}$ and we are done, or (2) u_{2l+1} is local-minimal. In this case $F \Delta \{u_1, \ldots, u_{2l+1}\} \in \mathcal{F}$. Since $u_1, \ldots, u_{2l}$ is, by induction hypothesis, a F-augmenting path then $u_1, \ldots, u_{2l}, u_{2l+1}$ is also F-augmenting and by Lemma D.1, there exists a subpath of $u_1, \ldots, u_{2l}, u_{2l+1}$ (possibly itself) which is F-local, in contradiction with the local-minimality of $u_1, \ldots, u_{2m+1}$. For paths of length $2m+2$ the proof is similar. Let F be in $\mathcal{F}$ and let $u_1, \ldots, u_{2m+2}$ be a F-local-minimal path. Since $u_1, \ldots, u_{2m}$ is also a F-local-minimal path, it suffices to show that $F \Delta \{u_1, \ldots, u_{2m+2}\} \in \mathcal{F}$. We will prove it by showing by induction on l, that for every $0 \leq l \leq m$, $F \Delta \{u_1, \ldots, u_{2l}, u_{2m+1}, u_{2m+2}\} \in \mathcal{F}$. The statement clearly holds for $l = 0$. Assume now that it is true for $l < m$. Since, $F \Delta \{u_1, \ldots, u_{2l}\}$, $F \Delta \{u_1, \ldots, u_{2l+2}\}$, and $F \Delta \{u_1, \ldots, u_{2l}, u_{2m+1}, u_{2m+2}\}$ are in $\mathcal{F}$, we can conclude that $u_{2l+1}, u_{2l+2}, u_{2m+1}, u_{2m+2}$ is a $F \Delta \{u_1, \ldots, u_{2l}\}$-local path. Thus, one of the following paths must be also local-minimal: (1) $u_{2l+1}, u_{2l+2}, u_{2m+1}, u_{2m+2}$, (2) u_{2l+1}, (3) $u_{2l+1}, v_{2l+2}, v_{2m+1}$, or (4) u_{2l+1}, v_{2m+2}. In case (1) $F \Delta \{u_1, \ldots, u_{2l+2}, u_{2m+1}, u_{2m+2}\}$ is in $\mathcal{F}$ and we are done. In case (2), $F \Delta \{u_1, \ldots, u_{2l+1}\} \in \mathcal{F}$. Thus, $u_1, \ldots, u_{2l+1}$ is a F-augmenting path. In consequence, by Lemma D.1 we can infer that there exists a F-local subpath of $u_1, \ldots, u_{2l+1}$, (and hence of $u_1, \ldots, u_{2m+2}$), in contradiction with the fact that $u_1, \ldots, u_{2m+2}$ is local-minimal. In case (3) we have that $F \Delta \{u_1, \ldots, u_{2l+2}, u_{2m+1}\}$ is in $\mathcal{F}$. In consequence, since by induction hypothesis $u_1, \ldots, u_{2l}$ is a F-augmenting path we have that $u_1, \ldots, u_{2l}, u_{2m+1}$ is also a F-augmenting path. By Lemma D.1 there exists a subpath $u_{k_1}, \ldots, u_{k_j}$ of $u_1, \ldots, u_{2l}, u_{2m+1}$ which is also F-local. We have two possibilities, if $k_j < 2m+1$, then $u_{k_1}, \ldots, u_{k_j}$ is also a F-local subpath of $u_1, \ldots, u_{2m+2}$, a contradiction. Otherwise, $F \Delta \{u_{2m+1}\} \in \mathcal{F}$ and consequently, $u_1, \ldots, u_{2m+1}$ is is a F-local subpath of $u_1, \ldots, u_{2m+2}$, a contradiction again. Finally, if (4) then $F \Delta \{u_1, \ldots, u_{2l+1}, u_{2m+2}\} \in \mathcal{F}$. Thus, $u_1, \ldots, u_{2l+1}, u_{2m+2}$ is a F-augmenting path and in consequence, there exists a subpath of $u_1, \ldots, u_{2l+1}, u_{2m+2}$ which is also F-local and a subpath of $u_1, \ldots, u_{2m+2}$, a contradiction. □

As a consequence of this result we can infer that the class $\mathcal{SL}$ of local Δ-matroids has an obstruction set of finite cardinality. More precisely, let $\mathcal{O}$ be the set of all non-local Δ-matroids $(E, \mathcal{F})$ such that $|E| \leq 4$, up to isomorphism. The following reasoning shows that $\mathcal{O}$ is an obstruction set for $\mathcal{SL}$:

Toward a contradiction, let $\mathcal{S}$ be the set of all non-local Δ-matroids that do not have any forbidden minor in $\mathcal{O}$ and select a Δ-matroid $(E, \mathcal{F})$ in $\mathcal{S}$ that does not have any minor in $\mathcal{S}$. Since $(E, \mathcal{F})$ is not local then it is not 2-local. Thus, there exists some $F \in \mathcal{F}$ and some F-local-minimal path $u_1, \ldots, u_n, n \leq 4$ that is not F-augmenting. Notice that $|E| > 4$ and consequently, there exists some $u \in E \setminus \{u_1, \ldots, u_n\}$. We study two possibilities: if $u \in F$, then let $(E', \mathcal{F}')$ be the elementary minor $(E, \mathcal{F})_u$, otherwise let $(E', \mathcal{F}')$ be the elementary minor $(E, \mathcal{F})_{\overline{u}}$. It is immediate to verify that $F \setminus \{u\}$ is a feasible set of $\mathcal{F}'$ and that $u_1, \ldots, u_n$ is a $F \setminus \{u\}$-local minimal non $F \setminus \{u\}$-augmenting. In consequence $(E', \mathcal{F}')$ is not local. Observe that since every minor of $(E', \mathcal{F}')$ is also a minor of $(E, \mathcal{F})$, $(E', \mathcal{F}')$ cannot have any forbidden minor in $\mathcal{O}$. In consequence $(E', \mathcal{F}') \in \mathcal{S}$, in contradiction with that fact that $(E, \mathcal{F})$ does not have any minor in $\mathcal{S}$.

Furthermore, it is then possible to carry out an exhaustive case analysis over the set of non-local Δ-matroids whose universe has at most 4 elements in order to obtain the minimal obstruction set of $\mathcal{SL}$.

Proposition E.3. *$\{L_1, L_2, L_3, L_4, L_5\}$ is the minimal obstruction set (up to isomorphism) of $\mathcal{SL}$ where $L_1 = (\{1,2,3\}, \{\emptyset, 2, 3, 12, 13\})$, $L_2 = (\{1,2,3\}, \{\emptyset, 3, 12, 13, 23\})$, $L_3 = (\{1,2,3\}, \{\emptyset, 2, 3, 12, 13, 23\})$, $L_4 = (\{1,2,3,4\}, \{\emptyset, 12, 13, 24, 34\})$, and $L_5 = (\{1,2,3,4\}, \{\emptyset, 12, 13, 23, 24, 34\})$.*

Proof. We start by noticing that every Δ-matroid of universe size at most 2 is local. Consequently, we have to consider only Δ-matroids of universe sizes 3 and 4.

Size 3. Let $(\{1,2,3\}, \mathcal{F})$ be a non-local Δ-matroid. Thus, there exists some $F \in \mathcal{F}$ and some F-local minimal path $u_1, \ldots, u_n$ that is not F-augmenting. Since every path of length at most 2 is also augmenting we can assume $n = 3$. Furthermore, by isomorphism, we can assume that $F = \emptyset$ and $u_1 = 1$, $u_2 = 2$, and $u_3 = 3$. Thus, we have that $\emptyset, 12, 3 \in \mathcal{F}$. Since $1, 2, 3$ is $\emptyset$-minimal, we have that $1 \notin \mathcal{F}$. Since $1, 2, 3$ is not $\emptyset$-augmenting we have that $123 \notin \mathcal{F}$. We also can conclude that $13 \in \mathcal{F}$ using the two-step axiom that defines Δ-matroids: observe that 13 is an step from 3 to 12, since none of the steps 1, 123 from 13 to 12 belongs to $\mathcal{F}$, we have that 13 must belong to $\mathcal{F}$. There are still two subsets, namely 2 and 23, for which still membership or no membership to $\mathcal{F}$ has to be established. It is easy to see that at least one of them must belong to $\mathcal{F}$: 23 is an step from 3 to 12, and 2 and 123 are the steps from 23 to 12. Since 123 is not in $\mathcal{F}$ then 2 or 23 must belong to $\mathcal{F}$. This gives three possibilities: $L_1 = \{\{1,2,3\}, \{\emptyset, 2, 3, 12, 13\}$, $L_2 = \{\emptyset, 3, 12, 13, 23\}$, and $L_3 = \{\emptyset, 2, 3, 12, 13, 23\}$. It is easy to check that L_1, L_2 and L_3 are Δ-matroids and that they are pairwise non-isomorphic. Thus, L_1, L_2 and L_3 are the minimal Δ-matroids of size 3 (up to isomorphism).

Size 4. Let $(\{1,2,3,4\}, \mathcal{F})$ be a non-local Δ-matroid. Thus, there exists some $F \in \mathcal{F}$ and some F-local minimal path $u_1, \ldots, u_n$ that is not F-augmenting. First notice that if $n \leq 3$ then $(\{1,2,3,4\}, \mathcal{F})$ has a forbidden minor and hence it is not minimal. Thus we can assume that $n = 4$. By isomorphism we can also assume that $F = \emptyset$, $u_1 = 1$, $u_2 = 2$, $y_3 = 3$, and $u_4 = 4$. Thus, we can infer that $\emptyset, 12, 34$ belong to $\mathcal{F}$. Since $1, 2, 3, 4$ is minimal we have that $1, 3$, and 14 do not belong to $\mathcal{F}$. Furthermore, since $1, 2, 3, 4$ is not augmenting we have that $1234 \notin \mathcal{F}$.

We can also conclude that $123 \notin \mathcal{F}$, since otherwise $2, 1, 3$ would be a 12-local minimal not 12-augmenting path, in contradiction with the fact that $(\{1,2,3,4\}, \mathcal{F})$ is minimal. Similarly, we can conclude that 134 is not in $\mathcal{F}$, since otherwise $4, 3, 1$ would be a 34-local minimal not 34-augmenting path.

Also, using the two-step axiom applied to 34 and 12 we can infer that 13 belongs to $\mathcal{F}$. First notice that 134 is an step from 34 to 12. Since 134 does not belong to $\mathcal{F}$ then at least one step from 134 to 12 must belong to $\mathcal{F}$. The steps from 134 to 12 are: 1234, 14, and 13. Since 1234 and 14 do not belong to $\mathcal{F}$, then we have that 13 belongs to $\mathcal{F}$.

Then we can conclude that $2 \notin \mathcal{F}$, otherwise $1, 3, 2$ would be a $\emptyset$-local minimal path not $\emptyset$-augmenting, in contradiction with the minimality of $(\{1,2,3,4\}, \mathcal{F})$.

Notice that if $4 \in \mathcal{F}$ then we can reach a contradiction in the following way. First, we have that 14 is an step from 4 to 13. Since 14 is not in $\mathcal{F}$ at least one of the steps from 14 to 13 must be in $\mathcal{F}$. The steps from 14 to 13 are 1 and 134 and none of them is in $\mathcal{F}$. Thus, we have just proved that $4 \notin \mathcal{F}$.

Then, we can conclude that 124 is not in $\mathcal{F}$. Otherwise, $2, 1, 4$ would be a 12-local minimal path non 12-augmenting, in contradiction with the minimality of $(\{1,2,3,4\}, \mathcal{F})$

Then, we can conclude that 24 must be in $\mathcal{F}$. Notice that 124 is an step from 12 to 34. Since 124 is not in $\mathcal{F}$, then some of the steps from 124 to 34 must belong to $\mathcal{F}$. Such steps are 24, 14, and 1234. Since 14 and 1234 does not belong to $\mathcal{F}$, we can conclude that 24 belongs to $\mathcal{F}$.

Also we can conclude that 234 is not in $\mathcal{F}$, since otherwise $3, 4, 2$ would be a 34-local-minimal not 34-augmenting path, in contradiction with the minimality of $(\{1,2,3,4\}, \mathcal{F})$.

So far, we have established the membership or not membership to $\mathcal{F}$ of all subsets of $\{1, 2, 3, 4\}$ but 23. Actually both possibilities for 23 are valid and they are not isomorphic, as it can be directly seen. Thus, we obtain two minimal Δ-matroids of universe size 4 (up to isomorphism) $L_4 = (\{1, 2, 3, 4\}, \{\emptyset, 12, 13, 24, 34\}$ and $L_5 = \{\{1, 2, 3, 4\}, \{\emptyset, 12, 13, 23, 24, 34\}\}$ $\square$