



Generalized minimum 0-extension problem and discrete convexity

Martin Dvorak¹ · Vladimir Kolmogorov¹ 

Received: 18 November 2022 / Accepted: 19 January 2024 / Published online: 7 March 2024
© The Author(s) 2024

Abstract

Given a fixed finite metric space (V, μ) , the *minimum 0-extension problem*, denoted as $0\text{-Ext}[\mu]$, is equivalent to the following optimization problem: minimize function of the form $\min_{x \in V^n} \sum_i f_i(x_i) + \sum_{ij} c_{ij} \mu(x_i, x_j)$ where $f_i : V \rightarrow \mathbb{R}$ are functions given by $f_i(x_i) = \sum_{v \in V} c_{vi} \mu(x_i, v)$ and c_{ij}, c_{vi} are given nonnegative costs. The computational complexity of $0\text{-Ext}[\mu]$ has been recently established by Karzanov and by Hirai: if metric μ is *orientable modular* then $0\text{-Ext}[\mu]$ can be solved in polynomial time, otherwise $0\text{-Ext}[\mu]$ is NP-hard. To prove the tractability part, Hirai developed a theory of discrete convex functions on orientable modular graphs generalizing several known classes of functions in discrete convex analysis, such as L^\natural -convex functions. We consider a more general version of the problem in which unary functions $f_i(x_i)$ can additionally have terms of the form $c_{uv;i} \mu(x_i, \{u, v\})$ for $\{u, v\} \in F$, where set $F \subseteq \binom{V}{2}$ is fixed. We extend the complexity classification above by providing an explicit condition on (μ, F) for the problem to be tractable. In order to prove the tractability part, we generalize Hirai's theory and define a larger class of discrete convex functions. It covers, in particular, another well-known class of functions, namely submodular functions on an integer lattice. Finally, we improve the complexity of Hirai's algorithm for solving $0\text{-Ext}[\mu]$ on orientable modular graphs.

Keywords Minimum 0-extension problem · Discrete convexity · Submodularity · L -convexity · Orientable modular graphs

Mathematics Subject Classification 90C27 · 05C12

✉ Vladimir Kolmogorov
vnk@ist.ac.at

Martin Dvorak
martin.dvorak@ist.ac.at

¹ Institute of Science and Technology Austria, Klosterneuburg, Austria

1 Introduction

Consider a metric space (V, μ) where V is a finite set and μ is a nonnegative function $V \times V \rightarrow \mathbb{R}$ satisfying the axioms of a metric: $\mu(x, y) = 0 \Leftrightarrow x = y$, $\mu(x, y) = \mu(y, x)$, $\mu(x, y) + \mu(y, z) \geq \mu(x, z)$ for all $x, y, z \in V$. We study optimization problems of the following form:

$$\min_{x \in V^n} f(x), \quad f(x) = \sum_{i \in [n]} f_i(x_i) + \sum_{1 \leq i < j \leq n} c_{ij} \mu(x_i, x_j) \quad (1)$$

where weights c_{ij} are nonnegative. If unary terms $f_i : V \rightarrow \mathbb{R}$ are allowed to be arbitrary nonnegative functions then this is a well-studied *Metric Labeling Problem* [22]. Another important special case is when the unary terms are given by

$$f_i(x_i) = \sum_{v \in V} c_{vi} \mu(x_i, v) \quad (2)$$

with nonnegative weights c_{vi} . This is a classical facility location problem, known as *multifacility location problem* [33]. It can be interpreted as follows: we are going to locate n new facilities in V , where the facilities communicate with each other and communicate with existing facilities in V . The cost of the communication is proportional to the distance. The goal is to find a location of minimum total communication cost. The multifacility location problem is also equivalent to the *minimum 0-extension problem* formulated by Karzanov [20]. We denote $0\text{-Ext}[\mu]$ to be class of problems of the form (1), (2).

Optimization problems of the form above have applications in computer vision and related clustering problems in machine learning [4, 11, 14, 22]. $0\text{-Ext}[\mu]$ also includes a number of basic combinatorial optimization problems. For example, the multiway cut problem on k vertices can be obtained by setting (V, μ) to be the uniform metric on $|V| = k$ elements; it can be solved in polynomial time (via a maximum flow algorithm) if $k = 2$, and is NP-hard for $k \geq 3$.

We explore a generalization of $0\text{-Ext}[\mu]$ in which the unary terms are given by

$$f_i(x_i) = \sum_{v \in V} c_{vi} \mu(x_i, v) + \sum_{U \in F} c_{Ui} \mu(x_i, U). \quad (3)$$

Here F is a fixed set of subsets of V , c_{vi} , c_{Ui} are nonnegative weights, and $\mu(x_i, U) = \min_{v \in U} \mu(x_i, v)$. We refer to this generalization as $0\text{-Ext}[\mu, F]$. In the facility location interpretation, allowing terms of the form $c_{Ui} \mu(x_i, U)$ means that the i -th facility can be “served” by any of the facilities in U , and it can choose to communicate with the closest facility to minimize the communication cost.

Note that $0\text{-Ext}[\mu] = 0\text{-Ext}[\mu, \emptyset]$. Furthermore, $0\text{-Ext}[\mu, 2^V]$, where $2^V = \{U \mid U \subseteq V\}$ is the set of all subsets of V , is the *restricted Metric Labeling Problem* [8], which is equivalent to the Metric Labeling Problem [6, Section 5.2].

1.1 Complexity classifications

The computational complexity of $0\text{-Ext}[\mu]$ has been established in [15, 20]. The tractability criterion is based on the properties of graph $H_\mu = (V, E, w)$ defined as the minimal undirected weighted graph whose path metric equals μ . Clearly, we have

$$E = \left\{ xy \in \binom{V}{2} \mid \forall z \in V - \{x, y\} : \mu(x, y) < \mu(x, z) + \mu(z, y) \right\}$$

and w is the restriction of μ to E . For brevity, we usually denote the elements of $\binom{V}{2}$ as xy instead of $\{x, y\}$.

In order to state the classification of $0\text{-Ext}[\mu]$, we need to introduce a few definitions.

Orientable modular graphs Let us fix metric μ . For nodes $x, y \in V$ let $I(x, y) = I_\mu(x, y)$ be the *metric interval* of x, y , i.e. the set of points $z \in V$ satisfying $\mu(x, z) + \mu(z, y) = \mu(x, y)$. Metric μ is called *modular* if for every triplet $x, y, z \in V$ the intersection $I(x, y) \cap I(y, z) \cap I(x, z)$ is non-empty. (Points in this intersection are called *medians* of x, y, z .) We say that graph H is *modular* if $H = H_\mu$ for a modular metric μ .

Let o be an edge-orientation of graph H with the relation \rightarrow_o on $V \times V$. This orientation is called *admissible* for H if, for every 4-cycle (x_1, x_2, x_3, x_4) , condition $x_1 \rightarrow_o x_2$ implies $x_4 \rightarrow_o x_3$. H is called *orientable* if it has an admissible orientation.

Theorem 1 ([20]) *If H_μ is not orientable or not modular then $0\text{-Ext}[\mu]$ is NP-hard.*

Theorem 2 ([15]) *If H_μ is orientable modular then $0\text{-Ext}[\mu]$ can be solved in polynomial time.*¹

Our results We extend the classification above to problems $0\text{-Ext}[\mu, F]$ in which all subsets $U \in F$ have cardinality 2, i.e. $F \subseteq \binom{V}{2}$. To formulate the tractability criterion, we need to introduce some definitions. Let o be an orientation of (H, F) , i.e. each edge of H is assigned an orientation, and each element of F is assigned an orientation. We say that o is *admissible* for (H, F) if it is admissible for H and, for every $\{x, y\} \in F$ with $x \rightarrow_o y$, the following holds: if P is a shortest x - y path in H then all edges of P are oriented according to o . We say that H is *F-orientable modular* if it is orientable modular and (H, F) admits an admissible orientation o . We can now formulate the main result of this paper.

Theorem 3 *If H_μ is F-orientable modular then $0\text{-Ext}[\mu, F]$ can be solved in polynomial time. Otherwise $0\text{-Ext}[\mu, F]$ is NP-hard.*

To prove the tractability part, we define *L-convex functions on extended modular complexes*, and show that they can be minimized in polynomial time. This generalizes *L-convex functions on modular complexes* introduced by Hirai in [15, 16].

¹ In this result μ is implicitly assumed to be rational-valued, since μ is treated as part of the input. The same remark applies to later results on VCSPs.

1.2 Discrete convex analysis

As Hirai remarks, the approach in [15, 16] had been inspired by discrete convex analysis developed in particular in [13, 27, 28, 30, 32] and [12, Chapter VII]. This is a theory of convex functions on integer lattice \mathbb{Z}^n , with the goal of providing a unified framework for polynomially solvable combinatorial optimization problems including network flows, matroids, and submodular functions. Hirai's work extends this theory to more general graph structures, in particular to orientable modular graphs, and provides a unified framework for polynomially solvable minimum 0-extension problems and related multiflow problems.

We develop a yet another generalization. To illustrate the relation to previous work, consider two fundamental classes of functions on the integer lattice $V = [k] = \{1, 2, \dots, k\}$: *submodular functions* and L^\natural -*convex functions*. These are functions $f : [k]^n \rightarrow \mathbb{R}$ satisfying conditions (4) and (5), respectively²

$$f(x) + f(y) \geq f(x \wedge y) + f(x \vee y) \quad \forall x, y \in [k]^n \quad (4)$$

$$f(x) + f(y) \geq f\left(\left\lceil \frac{1}{2}(x+y) \right\rceil\right) + f\left(\left\lfloor \frac{1}{2}(x+y) \right\rfloor\right) \quad \forall x, y \in [k]^n \quad (5)$$

where all operations are applied componentwise.

If, for example, $f(x) = \sum_i f_i(x_i) + \sum_{ij} f_{ij}(x_j - x_i)$, then f is submodular if all functions f_{ij} are convex, and f is L^\natural -convex if all functions f_i and f_{ij} are convex. The class of submodular functions on $[k]$ is strictly larger than the class of L^\natural -functions. However, L^\natural -convex functions possess additional properties that allow more efficient minimization algorithms, such as the Steepest Descent Algorithm [24, 29, 31].

The theory developed in [15, 16] covers L^\natural -convex functions and several other function classes, such as bisubmodular functions, k -submodular functions [18], skew-bisubmodular functions [19], and strongly-tree submodular functions [23]. However, it excludes submodular functions on $[k]$ for $k \geq 3$, which is a fundamental class of functions in discrete convex analysis. This paper fills this gap by introducing a unified framework that includes all classes of functions mentioned above.

Algorithms for solving 0-Ext $[\mu]$ and 0-Ext $[\mu, F]$ The tractability of 0-Ext $[\mu]$ for orientable modular μ was proven in [16] as follows. Given an instance $f : V^n \rightarrow \mathbb{R}$, Hirai defines a different instance $f_\times^* : (V^*)^n \rightarrow \mathbb{R}$ with the same minimum, where $|V^*| = O(|V|^2)$. Function f_\times^* is then minimized using the *Steepest Descent Algorithm* (SDA). This is an iterative technique that at each step computes a minimizer of f_\times^* in a certain local neighborhood of the current iterate (by solving a linear programming relaxation). We refer to this technique as the *SDA* approach*.

We present an alternative algorithm (for tractable classes of 0-Ext $[\mu, F]$) that minimizes function f directly via a version of SDA that we call \diamond -SDA (“diamond-SDA”). Both approaches terminate after $O(|V|)$ steps. However, one step of SDA* is more expensive than one step of SDA: the LP problem involves up to $O(|V^*|) =$

² We use a common notation $\overline{\mathbb{R}} = \mathbb{R} \cup \{\infty\}$, $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$, $\overline{\mathbb{Z}} = \mathbb{Z} \cup \{\infty\}$ where ∞ is an infinity element treated as follows: $\infty \cdot 0 = 0$, $x < \infty$ ($x \in \mathbb{R}$), $\infty + x = \infty$ ($x \in \mathbb{R}$), $x \cdot \infty = \infty$ ($a \in \mathbb{R} : a > 0$). We also denote \mathbb{R}_+ , \mathbb{Q}_+ , \mathbb{Z}_+ to be the sets of nonnegative elements of \mathbb{R} , \mathbb{Q} , \mathbb{Z} , respectively.

$O(|V|^2)$ labels per node in the former approach compared to $O(|V|)$ labels in the latter approach. Thus, we improve the complexity of solving $0\text{-Ext}[\mu]$.

Orthogonal generalizations of the minimum 0-extension problem In [16] Hirai considered the minimum 0-extension problem on *swm-graphs* (that generalize orientable modular graphs), and defined *L-extendable functions on swm-graphs*. Minimizing *L-extendable functions* on *swm-graphs* is an NP-hard problem (unless the graph is orientable modular); however, these functions admit a *discrete relaxation* on an orientable modular graph (which is an *L-convex function*). The relaxation can be minimized in polynomial time and yields a partial optimal solution for the original function. Based on this, Chalopin, Chepoi, Hirai and Osajda [5] obtained a 2-approximation algorithm for the minimum 0-extension problem on *swm-graphs*. They also developed the theory of *swm-graphs*.

In [17] Hirai and Mizutani considered minimum 0-extension problem for *directed metrics*, and provided some partial results (including a dichotomy for directed metrics of a star graph).

1.3 Valued Constraint Satisfaction Problems (VCSPs)

Results of this paper can be naturally stated in the framework of *Valued Constraint Satisfaction Problems* (VCSPs). This framework is defined below.

Let us fix a finite set D called a *domain*. A *cost function over D of arity n* is a function of the form $f : D^n \rightarrow \overline{\mathbb{R}}$. It is called *finite-valued* if $f(x) < \infty$ for all $x \in D^n$. We denote $\text{dom } f = \{x \in D^n \mid f(x) < \infty\}$. A *(VCSP) language over D* is a (possibly infinite) set Φ of cost functions over D . Language Φ is called *finite-valued* if all functions $f \in \Phi$ are finite-valued.

A *VCSP instance \mathcal{I}* is a function $D^n \rightarrow \overline{\mathbb{R}}$ given by

$$f_{\mathcal{I}}(x) = \sum_{t \in T} f_t(x_{v(t,1)}, \dots, x_{v(t,n_t)}). \quad (6)$$

It is specified by a finite set of variables $[n]$, finite set of terms T , cost functions $f_t : D^{n_t} \rightarrow \overline{\mathbb{R}}$ of arity n_t and indices $v(t, k) \in [n]$ for $t \in T, k = 1, \dots, n_t$. A solution to \mathcal{I} is a labeling $x \in [n]^V$ that minimizes $f_{\mathcal{I}}(x)$. Instance \mathcal{I} is called a *Φ -instance* if all terms f_t belong to Φ . The set of all Φ -instances is denoted as $\text{VCSP}(\Phi)$. Language Φ with finite $|\Phi|$ is called *tractable* if instances in $\text{VCSP}(\Phi)$ can be solved in polynomial time, and *NP-hard* if $\text{VCSP}(\Phi)$ is NP-hard. If $|\Phi|$ is infinite then Φ is called *tractable* if every finite $\Phi' \subseteq \Phi$ is tractable, and *NP-hard* if there exists finite $\Phi' \subseteq \Phi$ which is NP-hard.

A key algorithmic tool in the VCSP theory is the *Basic Linear Programming* (BLP) relaxation of instance \mathcal{I} . We refer to [25] for the description of this relaxation. We say that BLP *solves instance \mathcal{I}* if this relaxation is tight, i.e. its optimal value equals $\min_{x \in D^n} f_{\mathcal{I}}(x)$.

The following results are known; we refer to Sect. 6 for the definition of a “binary symmetric fractional polymorphism”.

Theorem 4 ([25]) *Let Φ be a finite-valued language. Then BLP solves Φ if and only if Φ admits a binary symmetric fractional polymorphism. If the condition holds, then an optimal solution of an Φ -instance can be computed in polynomial time.*

Theorem 5 ([34]) *If a finite-valued language Φ does not satisfy the condition in Theorem 4 then Φ is NP-hard.*

Application to the minimum 0-extension problem Consider again a metric space (V, μ) and subset $F \subseteq \binom{V}{2}$. For a set $U \subseteq V$ let $\delta_U : V \rightarrow \{0, \infty\}$ be the indicator function of set U , with $\delta_U(v) = 0$ iff $v \in U$. For brevity, we write $\delta_{\{u_1, \dots, u_k\}}$ as $\delta_{u_1 \dots u_k}$. Clearly, the minimum 0-extension problems introduced earlier can be equivalently defined by the following languages over domain $D = V$:

$$\begin{aligned} 0\text{-Ext}[\mu] &= \{\mu\} \cup \{\delta_u : u \in V\} \\ 0\text{-Ext}[\mu, F] &= \{\mu\} \cup \{\delta_u : u \in V\} \cup \{\delta_U : U \in F\} \end{aligned}$$

Note that the existence of the dichotomy given in Theorem 3 follows from Theorems 4 and 5 (but not the specific criterion for tractability).³

1.4 Summary of contributions

As described in earlier sections, our first contribution is complexity classification of $0\text{-Ext}[\mu, F]$ for subsets $F \subseteq \binom{V}{2}$ with an explicit criterion for tractability, which is achieved by generalizing Hirai's theory of modular complexes to *extended modular complexes*. While some of the proofs are relatively straightforward extensions of the corresponding proofs in [15, 16], there are also a number of proofs where we use novel techniques. We already mentioned a new \diamond -SDA algorithm that improves the complexity of solving the standard 0-minimum extension problem. In order to analyze this algorithm, we introduce new binary operations Δ, ∇, \diamond for (extended) modular complexes and establish their properties. Another key technical component that we use is the notion of *f-extremality* that we introduce. We believe that these concepts deepen our understanding of orientable modular graphs.

As our last contribution, we prove that the BLP relaxation directly solves L -convex functions on extended modular complexes. Previously, this was shown to hold (for standard modular complexes) only assuming that $P \neq NP$ (see Section 6 in [15]).

The rest of the paper is organized as follows. Section 2 reviews Hirai's theory and defines *L-convex functions on modular complexes* and the SDA algorithm for minimizing them. Section 3 generalizes this to *L-convex functions on extended modular complexes* and presents \diamond -SDA algorithm. Both sections use the notion of *submodular functions on valuated modular semilattices*, which are formally defined in Sect. 4. All proofs missing in Sect. 3 are given in Sects. 5 and 6.1; this completes the proof of the tractability direction of Theorem 3. The NP-hardness direction of Theorem 3 is proven in Sect. 6.2. Section 7 concludes the paper with a list of open problems.

³ Theorems 4 and 5 are not directly applicable to $0\text{-Ext}[\mu]$ since $0\text{-Ext}[\mu]$ is not finite-valued. However, we can get a finite-valued language by replacing functions $\delta_u : V \rightarrow \{0, \infty\}$ with functions $\mu_u : V \rightarrow \mathbb{R}$ defined via $\mu_u(x) = \mu(x, u)$. It is not difficult to show that such transformation does not affect the complexity of $0\text{-Ext}[\mu]$. A similar remark applies to $0\text{-Ext}[\mu, F]$.

2 Background on orientable modular graphs

Notation for graphs If H is a simple undirected graph and o its edge orientation, then the pair (H, o) can be viewed as a simple directed graph. We will usually denote this graph as $\Gamma = (V_\Gamma, E_\Gamma, w_\Gamma)$. Γ is called *oriented modular*, or a *modular complex*, if it is an admissible orientation of an orientable modular graph [15, 21].

We let \rightarrow_Γ be the edge relation of Γ , i.e. condition $u \rightarrow_\Gamma v$ means that there is an edge from u to v in Γ . When Γ is clear from the context, we may omit subscript Γ and write V, E, w, \rightarrow , etc. A *path* (u_0, u_1, \dots, u_k) in a directed graph Γ is defined as a path in the undirected version of Γ , i.e. for each i we must have either $u_i \rightarrow u_{i+1}$ or $u_{i+1} \rightarrow u_i$. An x - y path in Γ is a path from $x \in V$ to $y \in V$. With some abuse of notation we sometimes view Γ as a set of its nodes, and write e.g. $v \in \Gamma$ to mean $v \in V$.

Orbits For an undirected graph $H = (V, E, w)$ edges $e, e' \in E$ are called *projective* if there is a sequence of edges (e_0, e_1, \dots, e_m) with $(e_0, e_m) = (e, e')$ such that e_i, e_{i+1} are vertex-disjoint and belong to a common 4-cycle of H . Clearly, projectivity is an equivalence relation on E . An equivalence class of this relation is called an *orbit* [21]. Edge weights $w : E \rightarrow \mathbb{R}_{>0}$ are called *orbit-invariant* if $w(e) = w(e')$ for any pair of edges e, e' in the same orbit (equivalently, for vertex-disjoint edges e, e' belonging to a common 4-cycle).

Theorem 6 ([1, 21]) Consider undirected graph $H = H_\mu = (V, E, w)$.

- (a) If μ is modular then w is orbit-invariant, and path P is shortest in H if and only if it is shortest in $(V, E, 1)$.
- (b) The following conditions are equivalent: (i) H is (orientable) modular; (ii) w is orbit-invariant and $(V, E, 1)$ is (orientable) modular.

Metric spaces For a weighted directed or undirected graph $G = (V, E, w)$ let μ_G and d_G (or simply μ and d , when G is clear) be its path metrics w.r.t. edge lengths w and 1, respectively. If graph G is directed then edge orientations are again ignored.

For a metric space (V, μ) , subset $U \subseteq V$ is called *convex* if $I(x, y) \subseteq U$ for every $x, y \in U$. Note, if G is an orientable (or oriented) modular graph then the definitions of the metric interval $I(x, y)$ and of convex sets coincide for metric spaces (V_G, μ_G) and (V_G, d_G) (by Theorem 6).

Posets A modular complex Γ known to be an acyclic graph [15, Lemma 2.3], and thus induces a partial order \leq on V . Partially ordered sets (posets) play a key role in the study of oriented modular graphs. Below we describe basic facts about posets and terminology that we use, mainly following [15, 16].

Consider poset \mathcal{L} with relation \leq . For elements p, q with $p \leq q$, the *interval* $[p, q]$ is the set $\{x \in \mathcal{L} \mid p \leq x \leq q\}$. A *chain from p to q of length k* is a sequence $u_0 < u_1 < \dots < u_k$ with $(u_0, u_k) = (p, q)$, where notation $a < b$ means that $a \leq b$ and $a \neq b$. The length $r[p, q]$ of interval $[p, q]$ is defined as the maximum length of a chain from p to q . If \mathcal{L} has the lowest element (denoted as 0) then the *rank* $r(a)$ of element $a \in \mathcal{L}$ is defined by $r(a) = r[0, a]$, and elements of rank 1 are called *atoms*.

Element q covers p if $p < q$ and there is no element u with $p < u < q$. The *Hasse diagram* of \mathcal{L} is a directed graph on \mathcal{L} with the set of edges $\{p \rightarrow q \mid q \text{ covers } p\}$, and the *covering graph* of \mathcal{L} is the corresponding undirected graph.

A pair $x, y \in \mathcal{L}$ is said to be *upper-bounded* (resp. *lower-bounded*) if x, y have a common upper bound (resp. common lower bound). The lowest common upper bound, if exists, is denoted by $x \vee y$ (the “join” of x, y), and the greatest common lower bound, if exists, is denoted by $x \wedge y$ (the “meet” of x, y). \mathcal{L} is called a (*meet*-) *semilattice* if every pair $x, y \in \mathcal{L}$ has a meet, and it is called a *lattice* if every pair $x, y \in \mathcal{L}$ has both a join and a meet. If \mathcal{L} is a semilattice and $x, y \in \mathcal{L}$ are upper-bounded then $x \vee y$ is known to exist.

A (positive) *valuation* of a semilattice \mathcal{L} is a function $v : \mathcal{L} \rightarrow \mathbb{R}$ satisfying

$$v(q) - v(p) > 0 \quad \forall p, q \in \mathcal{L} : p < q \quad (7a)$$

$$v(p) + v(q) = v(p \wedge q) + v(p \vee q) \quad \forall p, q \in \mathcal{L} : p, q \text{ upper-bounded} \quad (7b)$$

In particular, if \mathcal{L} is a lattice then (7b) should hold for all $p, q \in \mathcal{L}$. A semilattice with valuation v will be called a *valuated semilattice*. We will view the Hasse diagram of a valuated semilattice \mathcal{L} (and the corresponding covering graph) as weighted graphs, where the weight of $p \rightarrow q$ is given by $v(q) - v(p) > 0$.

A lattice \mathcal{L} is called *modular* if for every $x, y, z \in \mathcal{L}$ with $x \leq z$ there holds $x \vee (y \wedge z) = (x \vee y) \wedge z$. A semilattice \mathcal{L} is called *modular* [2] if for every $p \in \mathcal{L}$ poset $(\{x \in \mathcal{L} \mid x \leq p\}, \leq)$ is a modular lattice, and for every $x, y, z \in \mathcal{L}$ the join $x \vee y \vee z$ exists provided that $x \vee y, y \vee z, z \vee x$ exist. It is known that a lattice \mathcal{L} is modular if and only if its rank function $r(\cdot)$ is a valuation [3, Chapter III, Corollary 1]. Furthermore, a (semi)lattice is modular if and only if its covering graph is modular, see [35, Proposition 6.2.1]; [2, Theorem 5.4]. The Hasse diagram of a (valuated) modular semilattice is known to be oriented modular [15, page 13].

Boolean pairs From now on we fix a modular complex $\Gamma = (V, E, w)$. Graph B is called a *cube graph* if it is isomorphic to the Hasse diagram of the Boolean lattice $\{0, 1\}^k$ for some $k \geq 0$. A pair of vertices (p, q) of Γ is called a *Boolean pair* if Γ contains cube graph B as a subgraph so that p and q are respectively the source and the sink of B . Let \sqsubseteq be the following relation on V : $p \sqsubseteq q$ iff (p, q) is a Boolean pair in Γ . We have $p \sqsubseteq p$ for any $p \in V$, and condition $p \sqsubseteq q$ implies that $p \leq q$. Graph Γ is called *well-oriented* if the opposite implication holds, i.e. if relations \leq and \sqsubseteq are the same. Note that relation \sqsubseteq is not necessarily transitive. We write $p \sqsubset q$ to mean $p \sqsubseteq q$ and $p \neq q$. Let Γ^\sqsubseteq be the graph with nodes V and edges $\{p \rightarrow q \mid p, q \in V, p \sqsubset q\}$. Clearly, Γ is a subgraph of Γ^\sqsubseteq (ignoring edge weights).

For a vertex $p \in V$ define the following subsets of V :

$$\mathcal{L}_p^\uparrow(\Gamma) = \{q \in V \mid p \leq q\} \quad \mathcal{L}_p^\downarrow(\Gamma) = \{q \in V \mid q \leq p\} \quad (8a)$$

$$\mathcal{L}_p^\pm(\Gamma) = \{q \in V \mid p \sqsubseteq q\} \quad \mathcal{L}_p^\mp(\Gamma) = \{q \in V \mid q \sqsubseteq p\} \quad \mathcal{L}_p^\pm(\Gamma) = \mathcal{L}_p^+(\Gamma) \cup \mathcal{L}_p^-(\Gamma) \quad (8b)$$

When Γ is clear from the context, we will omit it for brevity (and if Γ is not clear, we may write \leq_Γ and \sqsubseteq_Γ instead of \leq and \sqsubseteq). We view $\mathcal{L}_p^\uparrow, \mathcal{L}_p^\pm$ as posets with relation

\preceq , and $\mathcal{L}_p^\downarrow, \mathcal{L}_p^-$ as posets with the reverse of relation \preceq . Note, if Γ is well-oriented then $\mathcal{L}_p^+ = \mathcal{L}_p^\uparrow$ and $\mathcal{L}_p^- = \mathcal{L}_p^\downarrow$ for every node p of Γ .

Lemma 7 [[15, Proposition 4.1, Theorem 4.2, Lemma 4.14]] *Let Γ be a modular complex.*

- (a) *If elements a, b are upper-bounded then $a \vee b$ exists. Similarly, if a, b are lower-bounded then $a \wedge b$ exists.*
- (b) *Consider elements p, q with $p \preceq q$. Then $\mathcal{L}_p^\uparrow, \mathcal{L}_p^\downarrow, \mathcal{L}_p^+, \mathcal{L}_p^-$ are modular semilattices, and $[p, q] = \mathcal{L}_p^\uparrow \cap \mathcal{L}_p^\downarrow$ is a modular lattice. Furthermore, these (semi)lattices are convex in Γ , and function $v(\cdot)$ defined via $v(a) = \mu_\Gamma(p, a)$ is a valid valuation of these (semi)lattices.*

We will always view $\mathcal{L}_p^\uparrow, \mathcal{L}_p^\downarrow, \mathcal{L}_p^+, \mathcal{L}_p^-$ as **valuated** semilattices, where the valuation is defined as in the lemma.

L-convex functions Next, we review the notion of an *L-convex function* on a modular complex Γ introduced by Hirai in [15, 16]. The definition involves the following steps.

- First, Hirai defines the notion of a *submodular function* on a *valuated modular semilattice* \mathcal{L} . These are functions $f : \mathcal{L} \rightarrow \overline{\mathbb{R}}$ satisfying certain linear inequalities. Formulating these inequalities is rather lengthy, and we defer it to Sect. 4.
- Second, Hirai defines *2-subdivision* of Γ as the directed weighted graph $\Gamma^* = (V^*, E^*, w^*)$ constructed as follows:
 - (i) set $V^* = \{[p, q] \mid p \sqsubseteq q\}$;
 - (ii) for each $[p, q], [p, q'] \in V^*$ with $q \rightarrow_\Gamma q'$ add edge $[p, q] \rightarrow [p, q']$ to Γ^* with weight $w(qq')$;
 - (iii) for each $[p', q], [p, q] \in V^*$ with $p' \rightarrow_\Gamma p$ add edge $[p, q] \rightarrow [p', q]$ to Γ^* with weight $w(p'p)$.

It can be seen that graph (V^*, E^*) is the Hasse diagram of the poset (V^*, \subseteq) (which is the definition used in [16]).⁴ Hirai proves that graph Γ^* is oriented modular ([15, Theorem 4.3]) and well-oriented ([15, Lemma 2.14]). Consequently, poset $\mathcal{L}_{[p,p]}^+(\Gamma^*) = \mathcal{L}_{[p,p]}^\uparrow(\Gamma^*)$ is a valuated modular semilattice for each $p \in V$. For brevity, this poset will be denoted as $\mathcal{L}_p^*(\Gamma)$, or simply as \mathcal{L}_p^* .

- Each function $f : V \rightarrow \overline{\mathbb{R}}$ is extended to a function $f^* : V^* \rightarrow \overline{\mathbb{R}}$ via $f^*([p, q]) = f(p) + f(q)$.
- Function $f : V \rightarrow \overline{\mathbb{R}}$ is now called *L-convex on Γ* if (i) subset $\text{dom } f \subseteq V$ is connected in Γ^\square , and (ii) for every $p \in V$, the restriction of f^* to \mathcal{L}_p^* is submodular on valuated modular semilattice \mathcal{L}_p^* .

Minimum 0-extension problem and *L-convex functions* Next, we describe the relation between problem $0\text{-Ext}[\mu]$ and *L-convex functions* on Γ (where μ is an orientable modular metric and Γ is an admissible orientation of H_μ).

⁴ Compared to [15, 16], we chose to scale the weights of graph Γ^* by a factor of 2. Such scaling will not affect later theorems.

If Γ, Γ' are two modular complexes, then their Cartesian product $\Gamma \times \Gamma'$ is defined as the directed graph with the vertex set $V_\Gamma \times V_{\Gamma'}$ and the following edge set: there is an edge $(p, p') \rightarrow (q, q')$ iff either (i) $p = q$ and $p' \rightarrow_{\Gamma'} q'$, or (ii) $p' = q'$ and $p \rightarrow_\Gamma q$. The weight of the edge is $w_{\Gamma'}(p'q')$ in the first case and $w_\Gamma(pq)$ in the second case. The n -fold Cartesian product $\Gamma \times \dots \times \Gamma$ is denoted as Γ^n .

A Cartesian product $\mathcal{L} \times \mathcal{L}'$ of two posets $\mathcal{L}, \mathcal{L}'$ is defined in a natural way (with $(p, p') \leq (q, q')$ iff $p \leq q$ and $p' \leq q'$). It is straightforward to verify from definitions that if $\mathcal{L}, \mathcal{L}'$ are modular semilattices then so is $\mathcal{L} \times \mathcal{L}'$. If $\mathcal{L}, \mathcal{L}'$ are valuated modular semilattices with valuations $v_\mathcal{L}, v_{\mathcal{L}'}$ then $\mathcal{L} \times \mathcal{L}'$ is also assumed to be valuated with the valuation $v_{\mathcal{L} \times \mathcal{L}'}(p, p') = v_\mathcal{L}(p) + v_{\mathcal{L}'}(p')$.

Lemma 8 ([15, Lemma 4.7]) *Consider modular complexes Γ, Γ' and element $(p, p') \in \Gamma \times \Gamma'$.*

- (a) *Graph $\Gamma \times \Gamma'$ is a modular complex (i.e. oriented modular).*
- (b) $\mathcal{L}_{(p, p')}^\sigma(\Gamma \times \Gamma') = \mathcal{L}_p^\sigma(\Gamma) \times \mathcal{L}_{p'}^\sigma(\Gamma')$ for $\sigma \in \{-, +\}$.

Theorem 9 ([15, Theorem 4.8 and Lemma 4.9]; [16, Lemmas 4.1, 4.2 and Proposition 4.4]) *Consider modular complexes Γ, Γ' and functions $f, f' : \Gamma \rightarrow \mathbb{R}$ and $\tilde{f} : \Gamma \times \Gamma' \rightarrow \mathbb{R}$.*

- (a) *If f, f' are L -convex on Γ then $f + f'$ and $c \cdot f$ for $c \in \mathbb{R}_+$ are also L -convex on Γ .*
- (b) *If f is L -convex on Γ and $\tilde{f}(p, p') = f(p)$ for $(p, p') \in \Gamma \times \Gamma'$ then \tilde{f} is L -convex on $\Gamma \times \Gamma'$.*
- (c) *If \tilde{f} is L -convex on $\Gamma \times \Gamma'$ and $f(p) = \tilde{f}(p, p')$ for fixed $p' \in \Gamma'$ then f is L -convex on Γ .*
- (d) *The indicator function $\delta_U : V \rightarrow \{0, \infty\}$ of a d_Γ -convex set U is L -convex on Γ .*
- (e) *Function $\mu_\Gamma : \Gamma \times \Gamma \rightarrow \mathbb{R}_+$ is L -convex on $\Gamma \times \Gamma$.*
- (f) *If f is L -convex on Γ then the restriction of f to $\mathcal{L}_p^\sigma(\Gamma)$ is submodular on $\mathcal{L}_p^\sigma(\Gamma)$ for every $p \in \Gamma$ and $\sigma \in \{-, +\}$.*

It follows that an instance of $0\text{-Ext}[\mu]$ can be reduced to the problem of minimizing function $f : \Gamma^n \rightarrow \mathbb{R}, n \geq 1$ such that (i) f is L -convex on Γ^n , and (ii) $f(x_1, \dots, x_n)$ can be written as a sum of unary and pairwise terms that are L -convex on Γ and $\Gamma \times \Gamma$, respectively. This minimization problem is considered next.

Minimizing L -convex functions A key property of L -convex functions is that local optimality implies global optimality.

Theorem 10 ([15, Lemma 2.3]) *Let $f : V \rightarrow \mathbb{R}$ be an L -convex function on a modular complex Γ . If p is a local minimizer of f on $\mathcal{L}_p^\pm(\Gamma)$, i.e. $f(p) = \min_{q \in \mathcal{L}_p^\pm(\Gamma)} f(q)$, then it is also a global minimizer of f , i.e. $f(p) = \min_{q \in V} f(q)$.*

This theorem implies that the *Steepest Descent Algorithm* (SDA) given below is guaranteed to produce a global minimizer of f after a finite number of steps.

Algorithm 1: Steepest Descent for minimizing $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$

```

1 pick arbitrary  $x \in \text{dom } f$ 
2 while true do
3   compute  $y^+ \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^+(\Gamma^n)\}$  and  $y^- \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^-(\Gamma^n)\}$ 
4   pick  $y \in \arg \min \{f(y) \mid y \in \{y^-, y^+\}\}$ 
5   if  $f(y) = f(x)$  then return  $x$ , otherwise update  $x := y$ 

```

By Lemma 8, for $x = (x_1, \dots, x_n)$ we have $\mathcal{L}_x^\sigma(\Gamma^n) = \mathcal{L}_{x_1}^\sigma(\Gamma) \times \dots \times \mathcal{L}_{x_n}^\sigma(\Gamma)$ for $\sigma \in \{-, +\}$. By Theorem 9(f), function f is submodular on $\mathcal{L}_x^\sigma(\Gamma^n)$. The result below thus implies that the two minimization problems in line 3 can be solved in polynomial time assuming that f is an instance of 0-Ext[μ].

Theorem 11 [15, Theorem 3.9] *Consider VCSP instance $f : \mathcal{L}_1 \times \dots \times \mathcal{L}_n \rightarrow \overline{\mathbb{R}}$. Suppose that $\mathcal{L}_1, \dots, \mathcal{L}_n$ are valuated modular semilattices and function f is submodular on $\mathcal{L} = \mathcal{L}_1 \times \dots \times \mathcal{L}_n$. Then BLP relaxation solves f .*

To show that 0-Ext[μ] can be solved in polynomial time, a few additional definitions and results are needed. Elements p, q of a modular complex Γ are said to be \diamond -neighbors if $p, q \in [a, b]$ for some a, b with $a \sqsubseteq b$. Equivalently, p, q are \diamond -neighbors if $p \wedge q, p \vee q$ exist and $p \wedge q \sqsubseteq p \vee q$. Let Γ^\diamond be an undirected unweighted graph on nodes V_Γ such that p, q are connected by an edge in Γ^\diamond if and only if p, q are \diamond -neighbors. This graph is called a *thickening* of Γ . The distance from p to q in Γ^\diamond will be denoted as $d_\Gamma^\diamond(p, q)$, or simply as $d^\diamond(p, q)$. These distances will give a bound on the number of steps of SDA.

Theorem 12 [16, Theorem 4.7, Lemma 2.18] *Suppose that modular complex Γ is well-oriented, and function $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ is L -convex on Γ^n . SDA terminates after at most $2 + \max_{i \in [n]} d_\Gamma^\diamond(x_i, \text{opt}_i(f))$ iterations where x is the initial vertex and $\text{opt}_i(f) = \{x_i^* \mid x^* \in \arg \min f\} \subseteq V_\Gamma$ is the set of minimizers of f projected to the i -th component.*

Theorem 13 [5, Proposition 6.10], [16, Lemma 2.14] *If Γ is a modular complex then graph Γ^* is well-oriented.*

Theorem 14 [16, Proposition 4.9], [15, Lemma 4.7] *If Γ is a modular complex and function $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ is L -convex on Γ^n then function $f_\times^* : (\Gamma^*)^n \rightarrow \overline{\mathbb{R}}$ defined via $f_\times^*([x_1, y_1], \dots, [x_n, y_n]) = f(x) + f(y)$ is L -convex on $(\Gamma^*)^n$.*

From the results above one can now conclude that any instance $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ of 0-Ext[μ] for an oriented modular metric μ can be solved in polynomial time. Indeed, apply SDA to minimize function $f_\times^* : (\Gamma^*)^n \rightarrow \overline{\mathbb{R}}$. It will produce a minimizer $([x_1, y_1], \dots, [x_n, y_n])$ after at most $\text{diam}((\Gamma^*)^\diamond) + 2$ iterations; both x and y are minimizers of f . Hirai gives the bound $\text{diam}((\Gamma^*)^\diamond) \leq |V|^2$ in [16, Theorem 5.7]. As pointed out by an anonymous reviewer, this bound can be improved to $\text{diam}((\Gamma^*)^\diamond) \leq 2 \cdot \text{diam}(\Gamma^\diamond) \leq 2 \cdot |V|$.

Note that in the earlier paper [15] Hirai proved that 0-Ext[μ] can be solved in weakly polynomial time via a different algorithm, namely the Steepest Descent Algorithm (applied to $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$) but combined with a cost-scaling approach.

Remark 1 Our terminology is slightly different from that of [16]. In particular, \diamond -neighbors were called Δ -neighbors in [16], and a path in Γ^\square was called a Δ' -path.

3 Extended modular complexes

To prove our main tractability result from Theorem 3, we will introduce the following definition.

Definition 15 Let Γ be a modular complex. Binary relation \sqsubseteq on $V = V_\Gamma$ is called *admissible* if it coarsens \leq (i.e. $p \sqsubseteq q$ implies $p \leq q$), $p \sqsubseteq p$ for every $p \in V$, $p \sqsubseteq q$ for every edge $p \rightarrow q$, and the following conditions hold.

- (15a) Suppose that $p \sqsubseteq q$, $a \leq b$ and $a, b \in [p, q]$. Then $a \sqsubseteq b$.
- (15b) Suppose that $p \sqsubseteq q_1$, $p \sqsubseteq q_2$ and $q_1 \vee q_2$ exists. Then $p \sqsubseteq q_1 \vee q_2$.
- (15c) Suppose that $p_1 \sqsubseteq q$, $p_2 \sqsubseteq q$ and $p_1 \wedge p_2$ exists. Then $p_1 \wedge p_2 \sqsubseteq q$.

A pair (Γ, \sqsubseteq) where Γ is a modular complex and \sqsubseteq is admissible will be called an *extended modular complex*. With some abuse of notation, we will use letter Γ for an extended modular complex, and treat it as an oriented modular graph when needed. Note that previously we used notation \sqsubseteq for Boolean pairs in Γ . From now on we will write $p \sqsubseteq^{\text{Bp}} q$ if (p, q) is a Boolean pair in Γ , and reserve notation \sqsubseteq (or \sqsubseteq_Γ) for the binary relation that comes with an extended modular complex Γ . As before, we write $p \sqsubset q$ to mean $p \sqsubseteq q$ and $p \neq q$.

Proposition 16 Let Γ be a modular complex. (a) Relations \sqsubseteq^{Bp} and \leq are admissible for Γ . (b) If relation \sqsubseteq is admissible for Γ then $p \sqsubseteq^{\text{Bp}} q$ implies $p \sqsubseteq q$.

This proposition shows that a modular complex is a special case of an extended modular complex (obtained by setting \sqsubseteq to be \sqsubseteq^{Bp}). Also, \sqsubseteq^{Bp} and \leq are respectively the coarsest and the finest admissible relations. Next, we will show that many of the results in Sect. 2 still hold if relation \sqsubseteq^{Bp} is replaced with an arbitrary admissible relation \sqsubseteq .

From now on we fix an extended modular complex Γ . We define graphs Γ^\square and Γ^\diamond as in Sect. 2, and for $p \in V$ define posets $(\mathcal{L}_p^+, \mathcal{L}_p^-, \mathcal{L}_p^\pm) = (\mathcal{L}_p^+(\Gamma), \mathcal{L}_p^-(\Gamma), \mathcal{L}_p^\pm(\Gamma))$ as in Eq. (8). Note, in all cases \sqsubseteq now has the new meaning (it is the relation that comes with Γ).

Lemma 17 Let Γ be an extended modular complex, and p be its element. Then $\mathcal{L}_p^+, \mathcal{L}_p^-$ are modular semilattices that are convex in Γ . Furthermore, function $v(\cdot)$ defined via $v(a) = \mu_\Gamma(p, a)$ is a valid valuation of these (semi)lattices.

We define a 2-subdivision of Γ exactly as in Sect. 2; it is the graph $\Gamma^* = (V^*, E^*, w^*)$ where $V^* = \{[p, q] : p, q \in \Gamma, p \sqsubseteq_\Gamma q\}$. Using condition (15a), it is straightforward to verify that (V^*, E^*) is the Hasse diagram of the poset (V^*, \subseteq) . We will show the following result.

Theorem 18 If Γ is an extended modular complex then (V^*, E^*, w^*) is a modular complex (i.e. an oriented modular graph).




Fig. 1 Different 2-subdivisions for the chain $a < b < c < d$

Recall that if Γ is a modular complex (equivalently, an extended modular complex with the relation $\sqsubseteq_{\Gamma} = \sqsubseteq_{\Gamma}^{Bp}$) then Γ^* is well-oriented, i.e. $\sqsubseteq_{\Gamma^*} = \sqsubseteq_{\Gamma^*}^{Bp}$. For extended modular complexes this is not necessarily the case, as e.g. for the example in Fig. 1c. We will treat Γ^* as an extended modular complex with the relation $\sqsubseteq_{\Gamma^*} = \preceq_{\Gamma^*}$, so that we have $\mathcal{L}_{[p,p]}^+(\Gamma^*) = \mathcal{L}_{[p,p]}^{\uparrow}(\Gamma^*)$ by construction. Let us set $\mathcal{L}_p^* = \mathcal{L}_p^*(\Gamma) = \mathcal{L}_{[p,p]}^+(\Gamma^*) = \mathcal{L}_{[p,p]}^{\uparrow}(\Gamma^*)$. We can now define L -convex functions exactly as in the previous section:

- Function $f : V \rightarrow \overline{\mathbb{R}}$ is called *L-convex on extended modular complex* Γ if (i) subset $\text{dom } f \subseteq V$ is connected in Γ^{\sqsubseteq} , and (ii) for every $p \in V$, the restriction of f^* to $\mathcal{L}_p^* = \mathcal{L}_{[p,p]}^{\uparrow}(\Gamma^*)$ is submodular on valuated modular semilattice \mathcal{L}_p^* .

Cartesian products If (Γ, \sqsubseteq) , (Γ', \sqsubseteq') are two extended modular complexes, then their Cartesian product $(\Gamma, \sqsubseteq) \times (\Gamma', \sqsubseteq')$ is defined as the pair $(\Gamma \times \Gamma', \sqsubseteq_{\times})$ where $(p, p') \sqsubseteq_{\times} (q, q')$ iff $p \sqsubseteq q$ and $p' \sqsubseteq' q'$.

Lemma 19 Consider extended modular complexes Γ, Γ' and element $(p, p') \in \Gamma \times \Gamma'$.

- $\Gamma \times \Gamma'$ is an extended modular complex.
- $\mathcal{L}_{(p,p')}^{\sigma}(\Gamma \times \Gamma') = \mathcal{L}_p^{\sigma}(\Gamma) \times \mathcal{L}_{p'}^{\sigma}(\Gamma')$ for $\sigma \in \{-, +\}$.

Theorem 20 Consider extended modular complexes Γ, Γ' and functions $f, f' : \Gamma \rightarrow \overline{\mathbb{R}}$ and $\tilde{f} : \Gamma \times \Gamma' \rightarrow \overline{\mathbb{R}}$.

- If f, f' are L -convex on Γ then $f + f'$ and $c \cdot f$ for $c \in \mathbb{R}_+$ are also L -convex on Γ .
- If f is L -convex on Γ and $\tilde{f}(p, p') = f(p)$ for $(p, p') \in \Gamma \times \Gamma'$ then \tilde{f} is L -convex on $\Gamma \times \Gamma'$.
- If \tilde{f} is L -convex on $\Gamma \times \Gamma'$ and $f(p) = \tilde{f}(p, p')$ for fixed $p' \in \Gamma'$ then f is L -convex on Γ .
- The indicator function $\delta_U : V \rightarrow \{0, \infty\}$ is L -convex on Γ in the following cases:
 - U is a d_{Γ} -convex set;
 - $U = \{p, q\}$ for elements p, q with $p \sqsubseteq q$.
- Function $\mu_{\Gamma} : \Gamma \times \Gamma \rightarrow \mathbb{R}_+$ is L -convex on $\Gamma \times \Gamma$.
- If f is L -convex on Γ then the restriction of f to $\mathcal{L}_p^{\sigma}(\Gamma)$ is submodular on $\mathcal{L}_p^{\sigma}(\Gamma)$ for every $p \in \Gamma$ and $\sigma \in \{-, +\}$.

Theorem 21 Let $f : V \rightarrow \overline{\mathbb{R}}$ be an L -convex function on an extended modular complex Γ . If p is a local minimizer of f on $\mathcal{L}_p^\pm(\Gamma)$, i.e. $f(p) = \min_{q \in \mathcal{L}_p^\pm(\Gamma)} f(q)$, then it is also a global minimizer of f , i.e. $f(p) = \min_{q \in V} f(q)$.

For extended modular complex Γ let Φ_Γ be the language over domain $D = V_\Gamma$ that consists of all functions $f : D^n \rightarrow \overline{\mathbb{R}}$ such that f is L -convex on Γ^n . The results above imply that any instance \mathcal{I} of Φ_Γ can be solved by the Steepest Descent Algorithm (Algorithm 1), and each subproblem in line 2 can be solved in polynomial time. However, we do not know whether the number of steps would be polynomially bounded. To get a polynomial bound, we introduce an alternative algorithm, which we denote \diamond -SDA.

Algorithm 2: \diamond -SDA for minimizing $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$

```

1 pick arbitrary  $x \in \text{dom } f$ 
2 while true do
3   compute  $x^- \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^-(\Gamma^n)\}$  and  $x^+ \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^+(\Gamma^n)\}$ 
4   compute  $x^\diamond \in \arg \min \{f(y) \mid y \in \mathcal{L}_{x^-}^+(\Gamma^n) \cap \mathcal{L}_{x^+}^-(\Gamma^n)\}$ 
5   if  $f(x^\diamond) = f(x)$  then return  $x$ , otherwise update  $x := x^\diamond$ 

```

From Theorems 11 and 20(f) we can conclude that points x^-, x^+, x^\diamond in lines 3 and 4 can be computed in polynomial time via the BLP relaxation. To see this for x^\diamond , observe that function $f + \delta_U$ for convex set $U = \mathcal{L}_{x^-}^+ \cap \mathcal{L}_{x^+}^-$ is L -convex on Γ by Theorem 20(a,d), and thus its restriction to $\mathcal{L}_{x^-}^+$ is submodular on $\mathcal{L}_{x^-}^+$ by Theorem 20(f). (Alternatively, the claim can be deduced from Theorem 23 given later.)

Theorem 22 Let Γ be an extended modular complex and $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ be an L -convex function on Γ^n . \diamond -SDA algorithm applied to function f terminates after generating exactly $1 + \max_{i \in [n]} d_\Gamma^\diamond(x_i, \text{opt}_i(f))$ distinct points, where x is the initial vertex and $\text{opt}_i(f)$ is as defined in Theorem 12.

Remark 2 Suppose that $\sqsubset_\Gamma = \preceq_\Gamma$, and the initial vertex x in \diamond -SDA is computed as follows: pick some $x_0 \in \text{dom } f$ and then set either $x \in \arg \min \{f(y) \mid y \in \mathcal{L}_{x_0}^-\}$ or $x \in \arg \min \{f(y) \mid y \in \mathcal{L}_{x_0}^+\}$. It can be seen that in that case \diamond -SDA becomes equivalent to SDA: we will have $(x^-, x^+, x^\diamond) = (x, x^+, x^+)$ at even steps and $(x^-, x^+, x^\diamond) = (x^-, x, x^-)$ at odd steps, or vice versa. Thus, Theorem 22 generalizes Theorem 12 in two ways: from modular complexes to extended modular complexes, and by allowing relations \sqsubset_Γ and \preceq_Γ to be distinct.

Note that the algorithm for $0\text{-Ext}[\mu]$ described in the previous section required applying SDA on 2-subdivision Γ^* . This blows up the size of the graph and the size of LPs that need to be solved at each iteration by an up to a quadratic factor. \diamond -SDA provides an alternative that avoids such blow-up.

Remark 3 Algorithm 2 can also be specialized for minimizing L^\natural -convex functions (i.e. when Γ a directed path on consecutive integers and $a \sqsubset b$ iff $b = a + 1$). In this case

$\mathcal{L}_x^-(\Gamma^n) = [x - \mathbf{1}, x]$, $\mathcal{L}_x^+(\Gamma^n) = [x, x + \mathbf{1}]$ and $\mathcal{L}_{x^-}^+(\Gamma^n) \cap \mathcal{L}_{x^+}^-(\Gamma^n) = U_1 \times \dots \times U_n$ where

$$U_i = \begin{cases} \{x_i - 1, x_i\} & \text{if } (x_i^-, x_i^+) = (x_i - 1, x_i) \\ \{x_i, x_i + 1\} & \text{if } (x_i^-, x_i^+) = (x_i, x_i + 1) \\ \{x_i\} & \text{otherwise} \end{cases}.$$

Note that the resulting algorithm is different from previously proposed versions of SDA [24, 29, 31]. Thus, \diamond -SDA adds to the toolbox for L^\square -convex minimization.

We can now show the tractability part of Theorem 3. Suppose that graph H_μ is F -orientable modular for a metric space (V, μ) and subset $F \subseteq \binom{V}{2}$. Choose an admissible orientation of (H_μ, F) , and let Γ be the corresponding extended modular complex with the relation $\sqsubseteq = \preceq$. Clearly, for any $\{x, y\} \in F$ we have either $x \preceq y$ or $y \preceq x$. From Theorem 20 we conclude that $0\text{-Ext}[\mu, F] \subseteq \Phi_\Gamma$, and so $0\text{-Ext}[\mu, F]$ can be solved in polynomial time by the \diamond -SDA algorithm.

More generally, this shows tractability of $\text{VCSP}(\Phi_\Gamma)$ for an extended modular complex Γ assuming that a feasible solution of any Φ_Γ -instance can be computed in polynomial time. Our last result shows that $\text{VCSP}(\Phi_\Gamma)$ is tractable even without this assumption.

Theorem 23 *If Γ is an extended modular complex then BLP relaxation solves Φ_Γ , and an optimal solution of any Φ_Γ -instance can be computed in polynomial time.*

Note that previously Hirai remarked that BLP directly solves $0\text{-Ext}[\mu]$ for orientable modular metrics only assuming that $\mathbb{P} \neq \mathbb{NP}$, as a consequence of VCSP dichotomy for finite-valued languages (see Section 6 in [15]). Theorem 23 now establishes this fact unconditionally.

All proofs are given in Sects. 5 and 6. For Lemmas 17, 19 and Theorems 18, 20, 21 we mostly follow the proofs of the corresponding claims in [5, 15, 16] (replacing properties of relation \sqsubseteq^{Bp} with the properties of an admissible relation \sqsubseteq). To analyze \diamond -SDA (Theorem 22), we introduce new concepts, such as binary operations \triangle , ∇ , \diamond and the notion of f -extremality (see Sects. 5.6 and 5.9). Theorem 23 does not have an analogue in [5, 15, 16].

4 Submodular functions on a valuated modular semilattice

Let \mathcal{L} be a valuated modular semilattice with valuation v . This section gives the definition of a *submodular function on \mathcal{L}* , and thus completes the definition of L -convex functions.

Let $\Gamma = (V, E, w)$ be the Hasse diagram of \mathcal{L} where edge $p \rightarrow q$ is assigned weight $v(q) - v(p) > 0$. As discussed in Sect. 2, graph Γ is oriented modular. Denote $\mu = \mu_\Gamma$ and $d = d_\Gamma$. Recall that \mathcal{L} is viewed as a metric space with the metric μ , and the definitions of the metric interval $I(p, q)$, convex sets, etc would be the same for (\mathcal{L}, μ) and (\mathcal{L}, d) .

For $p \leq q$ let us denote $v[p, q] = v(q) - v(p)$. The following facts are known.




Fig. 2 **a** Each point in $I(p, q)$ is assigned a coordinate in \mathbb{R}^2 . The convex hull of these coordinates ($\text{Conv } I(p, q)$) is in gray. Distinct points may have the same coordinates (as some of the points shown in the interior of the gray region), but the coordinates of points in $\mathcal{E}(p, q) = \{u_0, u_1, \dots, u_k\}$ are guaranteed to be unique. **b** Definition of $\{\theta_i\}_i$ and $p \vee_\theta q$. First, define points $\alpha_{-1}, \alpha_0, \dots, \alpha_k$ in \mathbb{R}^2 as follows: set $\alpha_{-1} = (\sqrt{2}/2, 0)$, $\alpha_k = (0, \sqrt{2}/2)$ (so that $\|\alpha_k - \alpha_{-1}\| = 1$), and for $i \in [k-1]$ let α_i be the intersection of segment $[\alpha_{-1}, \alpha_k]$ and the line that goes through the origin and is perpendicular to the line passing through points $v_{pq}(u_k)$ and $v_{pq}(u_{k+1})$. Then $\theta_i = \|\alpha_i - \alpha_{-1}\|$ and $p \wedge_\theta q = u_i$ for each $i \in [0, k]$ and $\theta \in (\theta_{i-1}, \theta_i)$. **c** Bounded pair (p, q) . **d** Antipodal pair (p, q)

Lemma 24 [15, Lemma 2.15] *The following holds for each $p, q \in \mathcal{L}$ with $s = p \wedge q$.*

- (a) $\mu(p, q) = \mu(s, p) + \mu(s, q) = v[s, p] + v[s, q]$.
- (b) *The metric interval $I(p, q)$ is equal to the set of elements u that is represented as $u = a \vee b$ for some $(a, b) \in [s, p] \times [s, q]$, where such a representation is unique, and (a, b) equals $(u \wedge p, u \wedge q)$.*
- (c) *For $u, u' \in I(p, q)$ there holds $u \wedge u' = (u \wedge u' \wedge p) \vee (u \wedge u' \wedge q)$; in particular $u \wedge u' \in I(p, q)$.*

The construction in [15] can be described as follows (see Fig. 2a for a conceptual diagram).

- (i) For $u \in I(p, q)$, let $v_{pq}(u)$ be the vector in \mathbb{R}_+^2 defined by

$$v_{pq}(u) = (v[s, u \wedge p], v[s, u \wedge q]) \quad \text{where } s = p \wedge q \quad (9)$$

- (ii) Let $\text{Conv } I(p, q)$ denote the convex hull of vectors $v_{pq}(u)$ for all $u \in I(p, q)$.
- (iii) Let $\mathcal{E}(p, q)$ be the set of elements u in $I(p, q)$ such that $v_{pq}(u)$ is a maximal extreme point of $\text{Conv } I(p, q)$. (This set is called “ (p, q) -envelope”). Note that $p, q \in \mathcal{E}(p, q)$. Hirai proves that elements of $\mathcal{E}(p, q)$ receive unique coordinates [15, Lemma 3.1]:

$$v_{pq}(u) \neq v_{pq}(u') \quad \forall u \in \mathcal{E}(p, q), u' \in I(p, q) - \{u\} \quad (10)$$

- (iv) For $\theta \in [0, 1]$ define vector $c_\theta = (1-\theta, \theta)$. For points $p, q \in \mathcal{L}$ let $p \vee_\theta q$ be the point $u \in I(p, q)$ that maximizes $\langle c_\theta, v_{pq}(u) \rangle$, assuming that the maximizer is unique. If a maximizer is not unique then instead set $p \vee_\theta q = \perp$ (“undefined”).
- (v) By the property in Eq. (10), there are only a finite number of values $\theta \in [0, 1]$ such that $p \vee_\theta q = \perp$ for some $p, q \in \mathcal{L}$. Let Θ be the set of values $\theta \in [0, 1]$ such that $p \vee_\theta q \neq \perp$ for all $p, q \in \mathcal{L}$, then $\Theta \subseteq [0, 1]$ has measure 1.

We are now ready to define submodular functions on \mathcal{L} . A function $f : \mathcal{L} \rightarrow \overline{\mathbb{R}}$ is called *submodular* if every $p, q \in \mathcal{L}$ satisfies

$$f(p) + f(q) \geq f(p \wedge q) + \int_{\theta \in \Theta} f(p \vee_{\theta} q) d\theta \quad (11a)$$

This is equivalent to the inequality

$$f(p) + f(q) \geq f(p \wedge q) + \sum_{i=0}^k (\theta_i - \theta_{i-1}) f(u_i) \quad (11b)$$

where u_0, \dots, u_k are the sorted elements of $\mathcal{E}(p, q)$ with $(u_0, u_k) = (p, q)$ and $(\theta_{-1}, \theta_k) = (0, 1)$,

$$\theta_i = \frac{v[s_i, u_i]}{v[s_i, u_i] + v[s_i, u_{i+1}]} \quad \text{where } s_i = u_i \wedge u_{i+1}$$

for $i \in [0, k-1]$. (See Fig. 2b for a geometric interpretation of values θ_i).

Hirai also gives a simplified characterization of submodularity. Pair (p, q) is called *bounded* if $p \vee q$ exists, implying $\mathcal{E}(p, q) = \{p, q, p \vee q\}$ (Fig. 2c). Pair (p, q) is called *antipodal* if $\mathcal{E}(p, q) = \{p, q\}$ (Fig. 2d). Equivalently, (p, q) is antipodal if and only if

$$\begin{aligned} v[a, p]v[b, q] &\geq v[p \wedge q, a]v[p \wedge q, b] \\ \forall(a, b) : p \geq a \geq p \wedge q \leq b \leq q, a \vee b \text{ exists} \end{aligned} \quad (12)$$

We say that (p, q) is *special* if it is either bounded or antipodal. For special pairs inequality (11) reduces to

$$\begin{aligned} f(p) + f(q) &\geq f(p \wedge q) + f(p \vee q) \\ \text{if } (p, q) \text{ bounded} \end{aligned} \quad (13a)$$

$$\begin{aligned} v[p \wedge q, q]f(p) + v[p \wedge q, p]f(q) &\geq (v[p \wedge q, p] + v[p \wedge q, q])f(p \wedge q) \\ \text{if } (p, q) \text{ antipodal} \end{aligned} \quad (13b)$$

These inequalities are called respectively *submodularity* and \wedge -*convexity* inequalities for (p, q) .

Theorem 25 [15, Theorem 3.5] *Function $f : \mathcal{L} \rightarrow \overline{\mathbb{R}}$ is submodular if and only if it satisfies*

- (1) $\mathcal{E}(p, q) \subseteq \text{dom } f$ for $p, q \in \text{dom } f$;
- (2) the submodularity inequality for every bounded pair (p, q) ;
- (3) the \wedge -convexity inequality for every antipodal pair (p, q) .

We observe that condition (1) can be strengthened further (though we will not use this observation). For elements $p, q, u, u' \in \mathcal{L}$ we write $(p, q) \triangleleft (u, u')$ if

$u, u' \in I(p, q) - \{p, q\}$ and point $\frac{1}{2}(v_{pq}(u) + v_{pq}(u'))$ lies strictly above the segment $[v_{pq}(p), v_{pq}(q)]$ in \mathbb{R}^2 . Note that all elements $u \in \mathcal{E}(p, q) - \{p, q\}$ satisfy $(p, q) \triangleleft (u, u)$, and pair (p, q) is antipodal if and only if there is no $u \in \mathcal{L}$ with $(p, q) \triangleleft (u, u)$. The following theorem is proven in [10, Appendix A].

Theorem 26 *Function $f : \mathcal{L} \rightarrow \overline{\mathbb{R}}$ satisfies condition (1) of Theorem 25 if and only if it satisfies the following condition:*

- (1') *Suppose that $p, q \in \text{dom } f$ and $\mathcal{E}(p, q) \neq \{p, q\}$. Let u^- and u^+ be the elements in $\mathcal{E}(p, q) - \{p, q\}$ that are closest to p and to q , respectively. Then there exists $t \in \text{dom } f$ such that either $(p, q) \triangleleft (u^-, t)$ or $(p, q) \triangleleft (u^+, t)$.*

5 Proofs

Throughout the proofs we usually denote d, μ to be the distance functions in an extended modular complex Γ , and d^\square and d^\diamond to be the distance functions in Γ^\square and in Γ^\diamond respectively. We write $xy \in \Gamma^\square$ to indicate that nodes x, y are neighbors in Γ^\square (equivalently, that either $x \sqsubset y$ or $y \sqsubset x$). Note that $xy \in \Gamma^\square$ if and only if $d^\square(x, y) = 1$.

A sequence (u_0, u_1, \dots, u_k) of nodes in Γ is called a *shortest subpath* if $\mu(u_0, u_k) = \mu(u_0, u_1) + \mu(u_1, u_2) + \dots + \mu(u_{k-1}, u_k)$. Recall that for a modular graph Γ this is equivalent to the condition $d(u_0, u_k) = d(u_0, u_1) + d(u_1, u_2) + \dots + d(u_{k-1}, u_k)$. We will often implicitly use the following fact.

Proposition 27 *Consider elements p, q in a modular complex Γ . (a) If $p \wedge q$ exists then $(p, p \wedge q, q)$ is a shortest subpath. Conversely, if $p \geq x \leq q$ and (p, x, q) is a shortest subpath then $x = p \wedge q$. (b) If $p \vee q$ exists then $(p, p \vee q, q)$ is a shortest subpath. Conversely, if $p \leq x \leq q$ and (p, x, q) is a shortest subpath then $x = p \vee q$.*

Proof To see the claim, combine Lemma 7 and Lemma 24(a). \square

We say that sequence (x_1, x_2, x_3, x_4) is an *isometric rectangle* if (x_{i-1}, x_i, x_{i+1}) are shortest subpaths for all $i \in [4]$, where $x_j = x_{j \bmod 4}$. Isometric rectangles have the following properties.

Proposition 28 *If (x_1, x_2, x_3, x_4) is an isometric rectangle in graph Γ in then $d(x_1, x_2) = d(x_3, x_4)$ and $d(x_1, x_4) = d(x_2, x_3)$. Furthermore, if graph Γ is modular then for any $y_1 \in I(x_1, x_4)$ there exists $y_2 = I(x_2, x_3)$ (namely, a median of y_1, x_2, x_3) such that sequences (x_1, x_2, y_2, y_1) and (y_1, y_2, x_3, x_4) are isometric rectangles.*

Proof The definition of an isometric rectangle implies that $d(x_1, x_2) + d(x_2, x_3) = d(x_4, x_1) + d(x_3, x_4)$ and $d(x_2, x_3) + d(x_3, x_4) = d(x_1, x_2) + d(x_4, x_1)$, which in turns implies that $d(x_1, x_2) = d(x_3, x_4)$ and $d(x_1, x_4) = d(x_2, x_3)$. Now suppose that $y_1 \in I(x_1, x_4)$ and y_2 is a median of y_1, x_2, x_3 . (x_1, y_1, x_4, x_3) and (y_1, y_2, x_3) are shortest subpaths, and thus so is (x_1, y_1, y_2, x_3) . By a symmetric argument, (x_4, y_1, y_2, x_2) is also a shortest subpaths. This implies the claim. \square

5.1 Preliminaries

First, we state some known facts about orientable modular graphs that will be needed later in the proofs.

Lemma 29 ([2, Proposition 1.7], see [35, Proposition 6.2.6, Chapter I]) *A connected graph Γ with distances $d = d_\Gamma$ is modular if and only if*

- (1) Γ is bipartite, and
- (2) for vertices p, q and neighbors p_1, p_2 of p with $d(p_1, q) = d(p_2, q) = d(p, q) - 1$, there exists a common neighbor p^* of p_1, p_2 with $d(p^*, q) = d(p, q) - 2$.

Convex sets and gated sets Consider metric space (V, μ) . Recall that subset $U \subseteq V$ is called *convex* if $I(x, y) \subseteq U$ for every $x, y \in U$. It is called *gated* if for every $p \in V$ there exists unique $p^* \in U$, called the *gate of p at U* , such that $\mu(p, q) = \mu(p, p^*) + \mu(p^*, q)$ holds for every $q \in U$. The gate p^* will be denoted as $\text{Pr}_U(p)$ (“projection of p onto U ”). Thus, Pr_U is a map $V \rightarrow U$.

Theorem 30 [9] *Let A and A' be gated subsets of (V, μ) and let $B := \text{Pr}_A(A')$ and $B' := \text{Pr}_{A'}(A)$.*

- (a) Pr_A and $\text{Pr}_{A'}$ induce isometries, inverse to each other, between B' and B .
- (b) For $p \in A$ and $p' \in A'$, the following conditions are equivalent:
 - (i) $\mu(p, p') = \mu(A, A')$.
 - (ii) $p = \text{Pr}_A(p')$ and $p' = \text{Pr}_{A'}(p)$.
- (c) B and B' are gated, and $\text{Pr}_B = \text{Pr}_A \circ \text{Pr}_{A'}$ and $\text{Pr}_{B'} = \text{Pr}_{A'} \circ \text{Pr}_A$.

Now let us consider a modular graph $\Gamma = (V, E, w)$. Such graph induces two natural metrics on V , namely $\mu = \mu_\Gamma$ and $d = d_\Gamma$ (shortest path metrics w.r.t. edge lengths w and 1, respectively). In the light of Theorem 6, the definitions of convex sets, gated sets, gates and maps Pr_U would be the same for both metric spaces $((V, \mu)$ and (V, d)). In addition, convex and gated sets for such metrics coincide.

Lemma 31 ([7], see [15, Lemma 2.9]) *Let Γ be a modular graph. For $U \subseteq V$, the following conditions are equivalent:*

- (1) U is convex.
- (2) U is gated.
- (3) $\Gamma[U]$ is connected and $I(x, y) \subseteq U$ holds for every $p, q \in U$ with $d_\Gamma(p, q) = 2$.

Next, we review properties of modular complexes. A path $(p_0, p_1, p_2, \dots, p_k)$ in a directed acyclic graph $\Gamma = (V, E, w)$ is said to be *ascending* if $p_0 \prec p_1 \prec \dots \prec p_k$.

Lemma 32 [15, Lemma 4.13] *Let Γ be a modular complex. For $p, q \in V$ with $p \preceq q$, a (p, q) -path P is shortest if and only if P is an ascending path from p to q . In particular, $I(p, q) = [p, q]$, any maximal chain in $[p, q]$ has the same length, and the rank r of $[p, q]$ is given by $r(a) = d(a, p)$.*

Since set $[p, q]$ is convex in Γ by Lemma 7, one can define projection $\text{Pr}_{[p, q]} : V \rightarrow [p, q]$. The lemma below describes some properties of this projection.

Lemma 33 [15, Lemma 4.15] *Let Γ be a modular complex. For elements p, q, p', q' with $p \leq q$ and $p' \leq q'$ define*

$$u = \text{Pr}_{[p, q]}(p') \quad v = \text{Pr}_{[p, q]}(q') \quad u' = \text{Pr}_{[p', q']}(p) \quad v' = \text{Pr}_{[p', q']}(q) \quad (14)$$

Then we have:

- (1) $u \leq v, u' \leq v', \text{Pr}_{[p', q']}([p, q]) = [u', v']$, and $\text{Pr}_{[p, q]}([p', q']) = [u, v]$.
- (2) $[u, v]$ is isomorphic to $[u', v']$ by map $w \mapsto \text{Pr}_{[p', q']}(w)$.

5.2 Properties of extended modular complexes

In this section Γ is always assumed to be an extended modular complex on nodes V with relation \sqsubseteq .

Lemma 34 *Suppose that $p \sqsubseteq q_1, p \sqsubseteq q_2, q_1 \neq q_2$ and q is a common neighbor of q_1, q_2 (implying that $d(q_1, q_2) = 2$). Then $p \sqsubseteq q$.*

Proof Modulo symmetry, three cases are possible.

- $q_1 \rightarrow q \rightarrow q_2$. Since $p \sqsubseteq q_2$, condition (15a) gives $p \sqsubseteq q$.
- $q_1 \rightarrow q \leftarrow q_2$. Then $q = q_1 \vee q_2$ by Lemma 7(a), and condition (15b) gives $p \sqsubseteq q$.
- $q_1 \leftarrow q \rightarrow q_2$. Then $q = q_1 \wedge q_2$ by Lemma 7(a), implying $p \leq q$ (since p lower-bounds q_1, q_2). Condition (15a) gives $p \sqsubseteq q$.

□

Lemma 35 *Consider elements p, q, a, b such that (p, a, b) and (a, b, q) are shortest subpaths and $p \sqsubseteq q$. Then $a \sqsubseteq b$.*

Proof We use induction on $d(p, q) + d(p, a) + d(q, b)$. First, assume that (a, p, q) is not a shortest subpath. Let p' be a median of a, p, q , then $p' \in [p, q] - \{p\}$ and so $p' \sqsubseteq q$. The induction hypothesis for p', q, a, b gives the claim. We can thus assume that (a, p, q) is a shortest subpath. By a symmetric argument we can also assume that (b, q, p) is a shortest subpath. Sequence (p, q, b, a) is thus an isometric rectangle, and so $d(p, q) = d(a, b)$ and $d(p, a) = d(q, b)$. We now consider 5 possible cases.

- $d(p, a) = d(q, b) = 0$. The claim then holds trivially.
- $d(p, a) = d(q, b) = 1, q \rightarrow b$. Then $p \leq b$ and (p, a, b) is a shortest subpath, and so $p \rightarrow a \leq b$ by Lemma 32. We have $a \vee q \leq b$ and (a, b, q) is a shortest subpath; this implies that $a \vee q = b$. Since $p \sqsubseteq a$ and $p \sqsubseteq q$, we obtain $p \sqsubseteq a \vee q = b$ by condition (15b), and thus $a \sqsubseteq b$ since $a \in [p, b]$.
- $d(p, a) = d(q, b) = 1, p \leftarrow a$. This case is symmetric to the previous one.
- $d(p, a) = d(q, b) = 1, p \rightarrow a, b \rightarrow q$. We will show that $a \leq b$; this will imply that (p, a, b, q) is a shortest subpath by Lemma 32, contradicting condition

$d(p, q) = d(a, b)$. If $d(p, q) = d(a, b) = 0$ then the claim is trivial. If $d(p, q) = d(a, b) = 1$ then (p, q, b, a) is a 4-cycle, and so $p \rightarrow q$ implies that $a \rightarrow b$. Suppose that $d(p, q) = d(a, b) \geq 2$. By Proposition 28, there exist $x \in I(p, q) - \{p, q\}$ and $y \in I(a, b)$ such that (p, x, y, a) and (q, x, y, b) are isometric rectangles. We have $p \sqsubseteq x \sqsubseteq q$, so by the induction hypothesis $a \sqsubseteq y \sqsubseteq b$, implying the claim.

- $d(p, a) = d(q, b) \geq 2$. By Proposition 28, there exist $p' \in I(p, a)$ and $q' \in I(q, b)$ such that (p, q, q', p') and (p', q', b, a) are isometric rectangles. The induction hypothesis for elements p, q, p', q' gives $p' \sqsubseteq q'$. The induction hypothesis for elements p', q', a, b gives $a \sqsubseteq b$. \square

Corollary 36 Suppose that $p \sqsubseteq q$ and $A' \subseteq V$ is a convex set in Γ . Let $p' = \text{Pr}_{A'}(p)$ and $q' = \text{Pr}_{A'}(q)$. Then $p' \sqsubseteq q'$.

Lemma 17 (restated). Let Γ be an extended modular complex, and p be its element. Then $\mathcal{L}_p^+, \mathcal{L}_p^-$ are modular semilattices that are convex in Γ . Furthermore, function $v(\cdot)$ defined via $v(a) = \mu_\Gamma(p, a)$ is a valid valuation of these (semi)lattices.

Proof By symmetry, it suffices to consider only \mathcal{L}_p^+ . Note that \mathcal{L}_p^+ is a subset of \mathcal{L}_p^\uparrow such that for any $a, b \in \mathcal{L}_p^+$ we have (i) $a \wedge b \in \mathcal{L}_p^+$ and (ii) $a \vee b \in \mathcal{L}_p^+$ assuming that $a \vee b$ exists (by Definition 15). Thus, all claims except convexity follow from the corresponding properties of \mathcal{L}_p^\uparrow (see Lemma 7). Let us show the convexity. In view of Lemma 31, it suffices to show that for any $a, b \in \mathcal{L}_p^+$ with $d(a, b) = 2$ we have $I(a, b) \subseteq \mathcal{L}_p^+$. This claim follows directly from Lemma 34. \square

In the next two results we use the same notation d for d_Γ and d_{Γ^*} (since they can be distinguished by the arguments). Similarly, we use μ for both μ_Γ and μ_{Γ^*} .

Lemma 37 Let Γ^* be the 2-subdivision of Γ on nodes V^* . Then

$$d([p, q], [p', q']) = d(p, p') + d(q, q') \quad \forall [p, q], [p', q'] \in V^*$$

Proof If $P = ([p_0, q_0], \dots, [p_k, q_k])$ is a path in Γ^* from $[p_0, q_0] = [p, q]$ to $[p_k, q_k] = [p', q']$ then the length of P in Γ^* equals $\sum_{i=0}^{k-1} (d(p_i, p_{i+1}) + d(q_i, q_{i+1})) \geq d(p, p') + d(q, q')$; hence (\geq) holds.

To show equality, we use induction on $D = d(p, p') + d(q, q')$. The base case $D = 0$ is trivial; suppose that $D \geq 1$. It suffices to show that one of the following holds:

- p has neighbor a in Γ such that $a \sqsubseteq q$ and $d(a, p') = d(p, p') - 1$.
- q has neighbor b in Γ such that $p \sqsubseteq b$ and $d(b, q') = d(q, q') - 1$.

Indeed, if (i) holds then $d([a, q], [p', q']) = d(a, p') + d(q, q') = D - 1$ by induction, and hence $d([p, q], [p', q']) \leq (D - 1) + 1 = D$, as required. Case (ii) is similar. Alternatively, it would also suffice to show symmetrical cases when p' or q' have appropriate neighbors.

Define u, v, u', v' by (14). Note that $p \leq u \leq v \leq q$ and $p' \leq u' \leq v' \leq q'$. Modulo symmetry, two cases are possible.

- $p \neq u$, implying $p < u$. Let $a \in [p, u] \subseteq [p, q]$ be an out-neighbor of p in Γ (i.e. $p \rightarrow a \leq u$), then $a \subseteq q$ by condition (15a), and (p, a, p') is a shortest subpath since (p, a, u) and (p, u, p') are shortest subpaths. Thus, case (i) holds.
- $(p, q, p', q') = (u, v, u', v')$. Then (p, q, q', p') is an isometric rectangle. By Proposition 28 there exist $a \in I(p, p')$ and $b \in I(q, q')$ such that (p, q, b, a) and (p', q', b, a) are isometric rectangles and $d(q, b) = d(p, a) = 1$. We have $a \subseteq b$ by Lemma 35. If $a \rightarrow p$ then Lemma 32 for elements $a \leq q$ gives $a \leq b \rightarrow q$. By Lemma 7(a), $p \wedge b$ exists. Since (p, q, b) is a shortest subpath, we have $b \notin I(p, q) = [p, q]$, therefore $p \not\leq b$ and $p \wedge b = a$. Condition (15c) gives $a \subseteq q$, and so case (i) holds. If $q \rightarrow b$ then by a symmetric argument we conclude that case (ii) holds. The last remaining case $p \rightarrow a, b \rightarrow q$ is impossible by Lemma 32 for elements $p \leq q$. \square

Theorem 18 (restated). *If Γ is an extended modular complex then (V^*, E^*, w^*) is a modular complex (i.e. an oriented modular graph).*

Proof Any 4-cycle in Γ^* is represented as $([p, q], [p, q'], [p', q'], [p', q])$ or $([p', q], [p', q'], [p, q'], [p, q])$ for some edges $p \rightarrow p', q \rightarrow q'$ in Γ , or $([p, x], [p, y], [p, z], [p, w])$ or $([x, p], [y, p], [z, p], [w, p])$ for 4-cycle (x, y, z, w) and vertex p in Γ . This immediately implies that the orientation of Γ^* is admissible and orientation and w^* is orbit-invariant.

To show that Γ^* is modular, we are going to verify that Γ^* satisfies the two conditions of Lemma 29. If $[p, q]$ and $[p', q']$ are joined by an edge, then $d_\Gamma(p, q)$ and $d_\Gamma(p', q')$ have different parity. This implies that Γ^* is bipartite.

We next verify condition (2) of Lemma 29. Take intervals $[p, q], [p', q'] \in V^*$, and denote $D_p = d(p, p')$, $D_q = d(q, q')$, $D = d([p, q], [p', q']) = D_p + D_q$. Suppose further that we are given two distinct neighbors $[p_1, q_1], [p_2, q_2]$ of $[p, q]$ with $d([p_1, q_1], [p', q']) = d([p_2, q_2], [p', q']) = D - 1$. Our goal is to show the existence of a common neighbor $[p^*, q^*]$ of $[p_1, q_1], [p_2, q_2]$ with $d([p^*, q^*], [p', q']) = D - 2$.

Modulo symmetry, two cases are possible.

- $p_1 = p = p_2$. Condition $d([p_i, q_i], [p', q']) = D - 1$ implies that $d(q_i, q') = D_q - 1$ for $i = 1, 2$. By Lemma 29(2) for Γ , there is a common neighbor q^* of q_1, q_2 with $d(q^*, q') = D_q - 2$. By Lemma 34, $p \subseteq q^*$. Thus, $[p, q^*]$ is a desired common neighbor of $[p, q_1], [p, q_2]$.
- $p_1 = p, q_2 = q$. For better readability let us denote $s = p_2, t = q_1$. To summarize, we know that (p, s, p') and (q, t, q') are shortest subpaths, $d(p, s) = d(q, t) = 1$, $s \subseteq q, p \subseteq t, p \subseteq q, p' \subseteq q'$. It suffices to show $s \subseteq t$; this will imply that $[s, t]$ is a desired common neighbor of $[p_1, q_1], [p_2, q_2]$. Modulo symmetry, three subcases are possible. Case 1: $s \rightarrow p, t \rightarrow q$. Then $s \subseteq q$ and $s \leq p \leq t \leq q$, so condition (15a) gives $s \subseteq t$. Case 2: $p \rightarrow s, t \rightarrow q$. By Lemma 33, $\text{Pr}_{[p, q]}([p', q'])$ is equal to interval $[a, b]$ for $a = \text{Pr}_{[p, q]}(p')$ and $b = \text{Pr}_{[p, q]}(q')$. Note that $s, t, a, b \in [p, q]$. Since (p, s, p') and (s, a, p') are shortest subpaths, so is (p, s, a, p') and thus $s \leq a$. Similarly $b \leq t$. Thus $p \leq s \leq a \leq b \leq t \leq q$ and $p \subseteq q$ imply $s \subseteq t$ (by condition (15a)), as desired. Case 3: $s \rightarrow p, q \rightarrow t$. Observe that $x \vee y, x \wedge y$ are defined and belong to $[s, t]$ for all $x, y \in [s, t]$ by Lemma 7. Consider set $\text{Pr}_{[s, t]}([p', q'])$, which is equal to $[u, v]$ for $u = \text{Pr}_{[s, t]}(p')$

and $v = \text{Pr}_{[s,t]}(q')$ (Lemma 33). We must have $p \not\leq u$ (implying $p \wedge u = s$); otherwise (s, p, u) , (s, u, p') and thus (s, p, u, p') , (s, p, p') would be shortest subpaths, contradicting the assumption that (p, s, p') is a shortest subpath and $p \neq s$. Similarly, $v \vee q = t$. Note that $u \sqsubseteq v$ by Corollary 36. Define $a = u \wedge q$. We claim that $a \sqsubseteq t$. Indeed, we have $a \sqsubseteq q$ since $s \leq a \leq q$ and $s \sqsubseteq q$. It now suffices to show that there exists b such that $a \sqsubseteq b$ and $b \vee q = t$; the claim will then follow by condition (15b). If $a = u$ then we can take $b = v$. Otherwise, if $a < u$, take b to be an out-neighbor of a in $[a, u]$ (i.e. $a \rightarrow b \leq u$); we have $b \notin [a, q]$ since $a = u \wedge q$, and hence $b \vee q = t$. We have $p \sqsubseteq t$, $a \sqsubseteq t$ and $p \wedge a = p \wedge u \wedge q = s \wedge q = s$, so condition (15c) gives $s \sqsubseteq t$.

□

By combining Lemma 37, Theorem 18 and Theorem 6(a) we obtain:

Corollary 38 *Let Γ^* be the 2-subdivision of Γ on nodes V^* . Then*

$$\mu([p, q], [p', q']) = \mu(p, p') + \mu(q, q') \quad \forall [p, q], [p', q'] \in V^*$$

Lemma 39 *Let $f : \Gamma \rightarrow \mathbb{R}$ be an L -convex function on an extended modular complex Γ . Then for every $p \in \Gamma$ the restrictions of f to \mathcal{L}_p^- and to \mathcal{L}_p^+ are submodular functions.*

Proof Let us define $\mathcal{L}_p^{*+} = \{[p, q] : p \sqsubseteq q\} \subseteq \mathcal{L}_p^*$. By Lemma 37, any vertex in any shortest path between $[p, q]$ and $[p, q']$ is of the form $[p, u]$. Hence \mathcal{L}_p^{*+} is convex in Γ and in \mathcal{L}_p^* . Therefore, submodularity of f^* on \mathcal{L}_p^* implies f^* is submodular on \mathcal{L}_p^{*+} (by [15, Lemma 3.7(4)]). Obviously \mathcal{L}_p^{*+} is isomorphic to \mathcal{L}_p^+ by $[p, q] \mapsto q$. By using relation $f(q) = f^*([p, q]) - f(p)$ ($q \in \mathcal{L}_p^+$), we see the submodularity of f on \mathcal{L}_p^+ . The proof of submodularity of f on \mathcal{L}_p^- is symmetric. □

5.3 Proof of Proposition 16

Proposition 16 (restated). *Let Γ be a modular complex. (a) Relations \sqsubseteq^{Bp} and \leq are admissible for Γ . (b) If relation \sqsubseteq is admissible for Γ then $p \sqsubseteq^{\text{Bp}} q$ implies $p \sqsubseteq q$.*

To prove this proposition, we will need the following result. A modular lattice \mathcal{L} is called *complemented* if the maximal element $1_{\mathcal{L}}$ is a join of atoms. (This is one possible characterization of complemented modular lattices, see [3, Chapter IV, Theorem 4.1]).

Proposition 40 [5, Proposition 6.5] *Consider elements p, q in a modular complex Γ with $p \leq q$. Then $p \sqsubseteq^{\text{Bp}} q$ if and only if $[p, q]$ is a complemented modular lattice.*

We now proceed with the proof of Proposition 16.

(a) Checking admissibility of \leq is straightforward. Clearly, if B is a cube graph and p, q are elements of B with $p \leq q$ then the subgraph of B induced by $[p, q]$

is also a cube graph; this implies that \sqsubseteq^{Bp} satisfies condition (15a). Let us show that condition (15b) holds (condition (15c) is symmetric). Suppose that $p \sqsubseteq^{\text{Bp}} a$, $p \sqsubseteq^{\text{Bp}} b$ and $a \vee b$ exists. Let a_1, \dots, a_k be the atoms of a and b_1, \dots, b_ℓ be the atoms of b , then $a = a_1 \vee \dots \vee a_k$ and $b = b_1 \vee \dots \vee b_\ell$ by Proposition 40. Thus, $a \vee b = a_1 \vee \dots \vee a_k \vee b_1 \vee \dots \vee b_\ell$, and so $a \vee b$ is a join of atoms of $[p, a \vee b]$. Proposition 40 gives that $p \sqsubseteq^{\text{Bp}} a \vee b$.

(b) We use induction on $d(p, q)$. Consider elements p, q with $p \sqsubseteq^{\text{Bp}} q$, $d(p, q) \geq 2$. Using Proposition 40, we conclude that $q = a \vee b$ for some $a, b \in [p, q] - \{q\}$. We have $p \sqsubseteq^{\text{Bp}} a$ and $p \sqsubseteq^{\text{Bp}} b$ by part (a), and so $p \sqsubseteq a$ and $p \sqsubseteq b$ by the induction hypothesis. Condition (15b) gives $p \sqsubseteq q$.

5.4 Properties of the Cartesian product and L -convexity of the metric function

In this section we prove several properties related to the Cartesian product of extended modular complexes $\Gamma \times \Gamma'$. We denote $(V_\Gamma, V_{\Gamma'}, V_{\Gamma \times \Gamma'}) = (V, V', V_\times)$ for brevity, and use similar notation for other objects (e.g. relations $\leq, \sqsubseteq, \rightarrow$, distances d, μ , etc). Recall that $(p, p') \sqsubseteq_\times (q, q')$ iff $p \sqsubseteq q$ and $p' \sqsubseteq' q'$.

Lemma 19 (restated). *Consider extended modular complexes Γ, Γ' and element $(p, p') \in \Gamma \times \Gamma'$.*

- (a) $\Gamma \times \Gamma'$ is an extended modular complex.
- (b) $\mathcal{L}_{(p, p')}^\sigma(\Gamma \times \Gamma') = \mathcal{L}_p^\sigma(\Gamma) \times \mathcal{L}_{p'}^\sigma(\Gamma')$ for $\sigma \in \{-, +\}$.

Proof (a) In the light of Lemma 8(a), it suffices to verify that relation \sqsubseteq_\times is admissible, i.e. satisfies the properties in Definition 15. We need to show the following:

- $(p, p') \sqsubseteq_\times (q, q')$ implies $(p, p') \leq_\times (q, q')$.
- $(p, p') \sqsubseteq_\times (p, p')$ for every $(p, p') \in V_\Gamma \times V_{\Gamma'}$.
- $(p, p') \sqsubseteq_\times (q, q')$ for every edge $(p, p') \rightarrow_\times (q, q')$.
- If $(p, p') \sqsubseteq_\times (q, q')$, $(a, a') \leq_\times (b, b')$ and $(a, a'), (b, b') \in [(p, p'), (q, q')]$, then $(a, a') \sqsubseteq_\times (b, b')$.
- If $(p, p') \sqsubseteq (q_1, q'_1)$, $(p, p') \sqsubseteq (q_2, q'_2)$ and $(q_1, q'_1) \vee (q_2, q'_2)$ exists, then $(p, p') \sqsubseteq_\times (q_1, q'_1) \vee (q_2, q'_2)$.

Checking each property is mechanical, and is omitted.

(b) From definitions, $\mathcal{L}_{(p, p')}^+(\Gamma \times \Gamma') = \{(q, q') : p \sqsubseteq q, p' \sqsubseteq' q'\} = \mathcal{L}_p^+(\Gamma) \times \mathcal{L}_{p'}^+(\Gamma')$ as sets. Furthermore, the partial order is the same in both cases, and so $\mathcal{L}_{(p, p')}^+(\Gamma \times \Gamma')$ and $\mathcal{L}_p^+(\Gamma) \times \mathcal{L}_{p'}^+(\Gamma')$ also equal as posets (which are modular semilattices by Lemma 17). Finally, both semilattices are assigned the same valuation, namely $v_\times(q, q') = \mu_\times((p, p'), (q, q')) = \mu(p, q) + \mu'(p', q')$. The case $\sigma = -$ is symmetric. \square

Lemma 41 *Consider extended modular complexes Γ, Γ' . Graphs $(\Gamma \times \Gamma')^*$ and $\Gamma^* \times \Gamma'^*$ are isomorphic with the isomorphism given by $[(p, p'), (q, q')] \mapsto ([p, q], [p', q'])$ for $p, q \in \Gamma$, $p', q' \in \Gamma'$ with $p \sqsubseteq q$, $p' \sqsubseteq' q'$. Consequently,*

$\mathcal{L}_{(p,p')}^*(\Gamma \times \Gamma')$ and $\mathcal{L}_p^*(\Gamma) \times \mathcal{L}_{p'}^*(\Gamma')$ are isomorphic valuated modular semilattices for any $(p, p') \in \Gamma \times \Gamma'$.

Proof Clearly, the mapping defined in the lemma is a bijection between the nodes of $(\Gamma \times \Gamma')^*$ and the nodes of $\Gamma^* \times \Gamma'^*$. Checking that this bijection preserves edges and edge weights is mechanical. \square

Lemma 42 Consider extended modular complexes Γ, Γ' and functions $f, f' : \Gamma \rightarrow \overline{\mathbb{R}}$ and $\tilde{f} : \Gamma \times \Gamma' \rightarrow \overline{\mathbb{R}}$ such that $\text{dom } f$ is connected in Γ^\square .

- (a) If f is L -convex on Γ and $\tilde{f}(p, p') = f(p)$ for $(p, p') \in \Gamma \times \Gamma'$ then \tilde{f} is L -convex on $\Gamma \times \Gamma'$.
- (b) If \tilde{f} is L -convex on $\Gamma \times \Gamma'$ and $f(p) = \tilde{f}(p, p')$ for fixed $p' \in \Gamma'$ then f is L -convex on Γ .

Proof Since $(\Gamma \times \Gamma')^*$ and $\Gamma^* \times \Gamma'^*$ are isomorphic, for a function $\tilde{f}^* : (\Gamma \times \Gamma')^* \rightarrow \overline{\mathbb{R}}$ we can define function $\tilde{f}_\times^* : \Gamma^* \times \Gamma'^* \rightarrow \overline{\mathbb{R}}$ via $\tilde{f}_\times^*([p, q], [p', q']) = \tilde{f}^*([(p, p'), (q, q')])$. Clearly, \tilde{f}^* is submodular on $\mathcal{L}_{(x,x')}^*(\Gamma \times \Gamma')$ if and only if \tilde{f}_\times^* is submodular on $\mathcal{L}_x^*(\Gamma) \times \mathcal{L}_{x'}^*(\Gamma')$.

(a) Checking connectivity of $\text{dom } \tilde{f}$ in Γ^\square is straightforward. We can write $\tilde{f}_\times^*([p, q], [p', q']) = \tilde{f}(p, p') + \tilde{f}(q, q') = f(p) + f(q) = f^*([p, q])$. L -convexity of f means that f^* is submodular on \mathcal{L}_x^* for any $x \in \Gamma$. Therefore, by [15, Lemma 3.7(2)], function \tilde{f}_\times^* is submodular on $\mathcal{L}_x^*(\Gamma) \times \mathcal{L}_{x'}^*(\Gamma')$ for any $(x, x') \in \Gamma \times \Gamma'$, and thus \tilde{f}^* is submodular on $\mathcal{L}_{(x,x')}^*(\Gamma \times \Gamma')$.

(b) $f^*([p, q]) = f(p) + f(q) = \tilde{f}(p, p') + \tilde{f}(q, p') = \tilde{f}^*([(p, p'), (q, p')]) = \tilde{f}_\times^*([p, q], [p', p'])$. L -convexity of \tilde{f} means that \tilde{f}_\times^* is submodular on $\mathcal{L}_x^*(\Gamma) \times \mathcal{L}_{p'}^*(\Gamma')$. Therefore, by [15, Lemma 3.7(3)], function f^* is submodular on $\mathcal{L}_x^*(\Gamma)$ for any $x \in \Gamma$. \square

Lemma 43 [15, Lemma 4.18] Let Λ be a modular complex. The distance function μ_Λ is submodular on $\mathcal{L}_{(a,b)}^\uparrow(\Lambda \times \Lambda) = \mathcal{L}_a^\uparrow(\Lambda) \times \mathcal{L}_b^\uparrow(\Lambda)$ for every $(a, b) \in \Lambda \times \Lambda$.

Lemma 44 Let Γ be an extended modular complex. (a) Function $\mu_\Gamma : \Gamma \times \Gamma \rightarrow \mathbb{R}_+$ is L -convex on $\Gamma \times \Gamma$. (b) For each $p \in \Gamma$, function $\mu_{\Gamma,p} : \Gamma \rightarrow \mathbb{R}_+$ defined via $\mu_{\Gamma,p}(x) = \mu(x, p)$ is L -convex on Γ .

Proof It suffices to prove (a); claim (b) will then follow by Lemma 42(b). To be consistent with the notation in Lemma 42, denote $\Gamma' = \Gamma$ and $\tilde{f} = \mu_\Gamma$. By Corollary 38, for each $[p, q], [p', q'] \in V^*$ we have $\tilde{f}_\times^*([p, q], [p', q']) = \tilde{f}^*([(p, p'), (q, q')]) = \mu_\Gamma(p, p') + \mu_\Gamma(q, q') = \mu_{\Gamma^*}([p, q], [p', q'])$, i.e. $\tilde{f}_\times^* = \mu_{\Gamma^*}$. Consider $(x, y) \in \Gamma \times \Gamma$. Lemma 43 for $\Lambda = \Gamma^*$ and $a = [x, x], b = [y, y]$ gives that $\tilde{f}_\times^* = \mu_{\Gamma^*}$ is submodular on $\mathcal{L}_{[x,x]}^\uparrow(\Gamma^*) \times \mathcal{L}_{[y,y]}^\uparrow(\Gamma^*) = \mathcal{L}_x^*(\Gamma) \times \mathcal{L}_y^*(\Gamma)$, and therefore \tilde{f}^* is submodular on $\mathcal{L}_{(x,y)}^*(\Gamma \times \Gamma)$. \square

5.5 Proof of Theorem 21: Local optimality implies global optimality

In this section we prove the following result.

Theorem 21 (equivalent formulation). *Let $f : V \rightarrow \overline{\mathbb{R}}$ be an L -convex function on an extended modular complex Γ . Consider element p with $\min_{q \in \Gamma} f(q) < f(p) < \infty$. There exists element u with $f(u) < f(p)$ and $pu \in \Gamma^\square$.*

We use the same argument as in [15] (slightly rearranged).

Lemma 45 [15, Lemma 4.20] *Let f be a submodular function on a modular semi-lattice \mathcal{L} . For $p, q \in \mathcal{L}$ and $\alpha \in \mathbb{R}$ with $f(p) \leq \alpha$, $f(q) < \alpha$, there exists sequence (u_0, u_1, \dots, u_k) with $(u_0, u_k) = (p, q)$ such that for each $i \in [k]$ we have $f(u_i) < \alpha$ and elements u_{i-1}, u_i are comparable.*

Lemma 46 *Let f be an L -convex function on an extended modular complex Γ . Consider triplet (x, y, z) and $\alpha \in \mathbb{R}$ with $x \sqsubseteq y \sqsubseteq z$, $x \not\sqsubseteq z$, $f(x) \leq \alpha$, $f(y) < \infty$, $f(z) < \alpha$. There exists sequence (u_0, u_1, \dots, u_k) with $(u_0, u_k) = (x, z)$ such that for each $i \in [k]$ we have $f(u_i) < \alpha$ and $u_{i-1} \sqsubseteq u_i$.*

Proof We denote the desired sequence as $P(x, y, z)$, if exists. We prove that $P(x, y, z)$ exists using induction on $d(x, z) + d(x, y)$. We know that function f^* is submodular on \mathcal{L}_y^* . Let us apply inequality (11b) to elements $p = [x, y]$ and $q = [y, z]$ (with $p \wedge q = [y, y]$). Since $f^*(p) + f^*(q) = f(x) + f^*(p \wedge q) + f(z)$, there must exist $[a, b] \in \mathcal{E}(p, q) - \{p, q\}$ with $f(a) + f(b) = f^*([a, b]) \leq f(x) + f(z)$. Since $[a, b] \in I(p, q)$, we must have $a \in [x, y]$, $b \in [y, z]$ by Lemma 37 (and also $a \sqsubseteq b$). Condition $x \not\sqsubseteq z$ implies that $(a, b) \neq (x, z)$. Therefore, three cases are possible.

- $a = x, y < b < z$. Then we have $f(b) \leq f(z) < \alpha$ and $x \sqsubseteq b \sqsubseteq z$. Thus, we can set $P(x, y, z) = (x, b, z)$.
- $x < a < y, b = z$. Then we have $f(a) \leq f(x) \leq \alpha$ and $x \sqsubseteq a \sqsubseteq z$. Thus, we can set $P(x, y, z) = P(x, a, z)$, where we used the induction hypothesis (note that $d(x, a) < d(x, y)$).
- $x < a \leq b < z$. Since $f(a) + f(b) \leq f(x) + f(z) < 2\alpha$, one of the following must hold: $f(a) < \alpha$. Then we set $P(x, y, z) = (x, a, z)$ if $a \sqsubseteq z$, and $P(x, y, z) = (x, P(a, y, z))$ otherwise. $f(b) < \alpha$. Then we set $P(x, y, z) = (x, b, z)$ if $x \sqsubseteq b$, and $P(x, y, z) = (P(x, y, b), z)$ otherwise.

□

Corollary 47 *Consider elements x, y, z in an extended modular complex Γ such that $xy, yz \in \Gamma^\square$, $xz \notin \Gamma^\square$, $f(x) \leq \alpha$, $f(y) < \infty$, $f(z) < \alpha$. Then there exists path (u_0, u_1, \dots, u_k) in Γ^\square with $(u_0, u_k) = (x, z)$ such that $f(u_i) < \alpha$ for all $i \in [k]$.*

Proof Modulo symmetry, two cases are possible:

- $x \sqsupset y \sqsubseteq z$. The claim then follows from Lemma 45, since function f is submodular on \mathcal{L}_y^+ by Lemma 39.
- $x \sqsubseteq y \sqsupset z$. The claim then follows from Lemma 46.

□

We now proceed with the proof of Theorem 21. For value $\alpha \in \mathbb{R}$ let Γ_α^\square and $\Gamma_{<\alpha}^\square$ be the subgraphs of Γ^\square induced by nodes p with $f(p) \leq \alpha$ and with $f(p) < \alpha$, respectively.

Lemma 48 Γ_α^\square is connected for any $\alpha \geq \alpha^* \stackrel{\text{def}}{=} \min_{p \in \Gamma} f(p)$.

Proof Suppose the claim is false. For a sufficiently large α graph Γ_α^\square is connected, since $\text{dom } f$ is connected in Γ^\square . Thus, there must exist $\alpha > \alpha^*$ such that Γ_α^\square is connected but $\Gamma_{<\alpha}^\square$ is disconnected. There must exist a pair of vertices p, p' belonging to different components of $\Gamma_{<\alpha}^\square$; in particular, $f(p) < \alpha$ and $f(p') < \alpha$, and there exists path $P = (p_0, p_1, \dots, p_k)$ in Γ^\square with $(p_0, p_k) = (p, p')$, $k \geq 2$ and $f(p_i) \leq \alpha$ for all $i \in [k-1]$. Pick such p, p', P so that $k \geq 2$ is minimum. The minimality of k implies that $p_{k-2}p_k \notin \Gamma^\square$. Apply Corollary 47 to elements $(x, y, z) = (p_{k-2}, p_{k-1}, p_k)$. We obtain a path (u_0, \dots, u_ℓ) in Γ^\square between $u_0 = p_{k-2}$ and $u_\ell = p_k$ with $f(u_i) < \alpha$ for all $i \in [\ell]$. Note that u_1, p' are in the same connected component of $\Gamma_{<\alpha}^\square$ (which is different from that of p), and path $(p_0, \dots, p_{k-2}, u_1)$ has shorter length compared to P . This contradicts the minimality of k . \square

Now consider element $p \in \Gamma$ which is not a global minimum, i.e. $f(p) > \alpha^*$. Connectivity of Γ_α^\square for $\alpha = f(p)$ implies that there exists path $P = (p_0, p_1, \dots, p_k)$ in Γ^\square with $p_0 = p$, $f(p_i) \leq \alpha$ for $i \in [k]$, and $f(p_k) < \alpha$. Let us pick such P so that k is minimum. It suffices to show that $k = 1$; clearly, this would imply Theorem 21. Suppose not, i.e. $k \geq 2$. The minimality of k implies that $p_{k-2}p_k \notin \Gamma^\square$. Applying Corollary 47 to elements $(x, y, z) = (p_{k-2}, p_{k-1}, p_k)$ gives element u_1 so that $f(u_1) < \alpha$ and $p_{k-2}u_1 \in \Gamma^\square$. This contradicts the minimality of k .

We can strengthen Theorem 21 as follows.

Lemma 49 Let $f : V \rightarrow \overline{\mathbb{R}}$ be an L -convex function on an extended modular complex Γ and p, q be distinct elements in $\text{dom } f$. (a) If $f(p) \geq f(q)$ then there exists element $u \in I(p, q)$ such that $f(u) \leq f(p)$ and $pu \in \Gamma^\square$. Additionally, if $f(p) > f(q)$ then $f(u) < f(p)$. (b) There exists element $u \in I(p, q) \cap \text{dom } f$ such that $pu \in \Gamma^\square$.

Proof (a) For element $x \in \Gamma$ define function $\mu_x : \Gamma \rightarrow \mathbb{R}$ via $\mu_x(u) = \mu(x, u)$. By Lemma 44, function μ_x is L -convex on Γ . Now for value $C \geq 0$ define function $f_C : \Gamma \rightarrow \overline{\mathbb{R}}$ via $f_C(u) = f(u) + C(\mu_p(u) + \mu_q(u) - \mu(p, q))$. It is straightforward to check that f_C is L -convex on Γ . (Note that $\text{dom } f_C = \text{dom } f$, since functions μ_p, μ_q are finite-valued). We have the following properties: (i) $f_C(u) = f(u)$ for all $u \in I(p, q)$; (ii) if C is sufficiently large then $f_C(u) > f(p)$ for all $u \in V_\Gamma - I(p, q)$. Applying Lemma 48 to function f_C gives the first claim, while applying Theorem 21 to function f_C gives the second claim.

(b) We can assume w.l.o.g. that $f(p) > f(q)$, since adding function of the form $C\mu_q$, $C \geq 0$ to f preserves L -convexity of f and does not affect the statement. The claim now follows from part (a). \square

5.6 Graph thickening and operations $\triangle, \nabla, \diamond$

Let us recall some definitions from Sects. 2, 3. Elements p, q of an extended modular complex Γ are said to be \diamond -neighbors if $p, q \in [a, b]$ for some a, b with $a \sqsubseteq b$. Equivalently, p, q are \diamond -neighbors if $a \wedge b, a \vee b$ exist and $a \wedge b \sqsubseteq a \vee b$. Let Γ^\diamond be an undirected unweighted graph on nodes V_Γ such that p, q are connected by an

edge in Γ^\diamond if and only if p, q are Δ -neighbors. This graph is called a *thickening* of Γ . A path in Γ^\diamond is called a \diamond -path. The shortest path distance between p and q in Γ^\diamond will be denoted as $d^\diamond(p, q)$. These concepts were introduced in [5] in the case when $\sqsubseteq = \sqsubseteq^{\text{Bp}}$; we now use them for an arbitrary admissible relation \sqsubseteq . In order to work with Γ^\diamond , in this section we introduce binary operations Δ, ∇, \diamond on an extended modular complex Γ and establish some of their properties. They will be used, in particular, for proving Theorem 22 (bound on the number of steps of \diamond -SDA), and for proving that the sum of L -convex functions is L -convex.

For vertices $p, q \in \Gamma$ let \leq_{pq} be a partial order on $I(p, q)$ defined as follows: $x \leq_{pq} y$ iff (p, x, y, q) is a shortest subpath. For vertices $x, y \in I(p, q)$ this is equivalent to the condition $x \in I(p, y)$, or to the condition $y \in I(x, q)$. Clearly, $(I(p, q), \leq_{pq})$ is a poset with the minimal element p and the maximal element q . We will need the following result.

Theorem 50 ([5, Theorem 6.1]) *Let Γ be a modular complex. For every $p, q \in \Gamma$ poset $(I(p, q), \leq_{pq})$ is a modular lattice.*

We denote the meet and join operations in this poset as \wedge_{pq} and \vee_{pq} , respectively. Clearly, for each $x, y \in I(p, q)$ the meet $x \wedge_{pq} y$ is a median of p, x, y , and in fact it is the unique median (since every median m of p, x, y belongs to $I(p, q)$ and satisfies $m \leq_{pq} x$ and $m \leq_{pq} y$). Similarly, $x \vee_{pq} y$ is the unique median of x, y, q .

For elements $p, q \in \Gamma$ we define $p \Delta q$ to be the gate of q at \mathcal{L}_p^+ (which is a convex subset of Γ by Lemma 17): $p \Delta q = \text{Pr}_{\mathcal{L}_p^+}(q)$. Similarly, we let $p \nabla q = \text{Pr}_{\mathcal{L}_p^-}(q)$. Clearly, by the definition of the gate and by Lemma 32 we have

$$I(p, q) \cap \mathcal{L}_p^+ = I(p, p \Delta q) = [p, p \Delta q] = \{u \mid (p, u, p \Delta q, q) \text{ is a shortest subpath}\} \quad (15a)$$

$$I(p, q) \cap \mathcal{L}_p^- = I(p, p \nabla q) = [p \nabla q, p] = \{u \mid (p, u, p \nabla q, q) \text{ is a shortest subpath}\} \quad (15b)$$

We also define $p \diamond q = (p \Delta q) \vee_{pq} (p \nabla q)$.

Lemma 51 *Consider elements p, q of an extended modular complex Γ . There holds $p \wedge (p \diamond q) = p \nabla q \sqsubseteq p \Delta q = p \vee (p \diamond q)$ (and consequently $p, p \diamond q$ are \diamond -neighbors). If $p \neq q$ then $p \neq p \diamond q$.*

Proof Denote $(a, b, c) = (p \nabla q, p \Delta q, p \diamond q)$. We know that $a \sqsubseteq p \sqsubseteq b$ and $(p, a, c), (p, b, c), (a, c, b)$ are shortest subpaths (since c is a median of a, b, q and $a, b \in I(p, q)$). From Lemma 35 we conclude that $a \sqsubseteq c \sqsubseteq b$. Since $p \geq a \leq c$ and (p, a, c) is a shortest subpath, we must have $a = p \wedge c$. Similarly, $b = p \vee c$. Condition (15b) gives $a \sqsubseteq b$. Now suppose that $p \neq q$, and consider shortest p - q path (p, u, \dots, q) . If $p \rightarrow u$ then $p \Delta q \neq p$ and $p \diamond q \neq p$, and if $p \leftarrow u$ then $p \nabla q \neq p$ and $p \diamond q \neq p$. \square

For elements p, q an integer $k \geq 0$ define $p \diamond^k q$ as follows: $p \diamond^0 q = p$, and $p \diamond^{k+1} q = (p \diamond^k q) \diamond q$ for $k \geq 0$. Clearly, $p \diamond^k q = q$ for some index $k \geq 0$. If k is

the minimum such index then the sequence $(p \diamond^0 q, p \diamond^1 q, \dots, p \diamond^k q)$ will be called the *normal p - q path*. By the previous lemma, it is a \diamond -path (i.e. a path in Γ^\diamond).

For node $p \in \Gamma$ and integer $k \geq 0$ denote $B_k^\diamond(p) = \{q \mid d^\diamond(p, q) \leq k\}$ (a ball of radius k in Γ^\diamond). We will show later that $B_k^\diamond(p)$ is convex in Γ for any p and k . First, we establish this for $k = 1$.

Lemma 52 *For each $p \in \Gamma$, set $B_1^\diamond(p)$ is convex in Γ .*

Proof By Lemma 31, it suffices to show the following: if x, x' are distinct \diamond -neighbors of p and y is a common neighbor of x, x' in Γ then p, y are \diamond -neighbors. Denote $(a, b) = (p \wedge x, p \vee x)$ and $(a', b') = (p \wedge x', p \vee x')$, then $a \sqsubseteq b$ and $a' \sqsubseteq b'$. Recall that sets $[a, b]$ and $[a', b']$ are convex in Γ by Lemma 7. If $d(p, y) = d(p, x) - 1$ then $y \in I(p, x)$, and so by convexity of $[a, b]$ we have $y \in [a, b]$, implying that p, y are \diamond -neighbors. We can thus assume that $d(p, y) = d(p, x) + 1$, and also that $d(p, y) = d(p, x') + 1$ (by a symmetric argument). Modulo symmetry, two cases are possible.

- $x \rightarrow y \rightarrow x'$. Since (p, a, x) and (p, x, y) are shortest subpaths, so is (p, a, y) . Since (p, b', x') and (p, x', y') are shortest subpaths, so is (p, b', y) . (a, p, b') and (a, y, b') are also shortest subpaths (since $a \leq p < b'$ and $a \leq y < b'$). Thus, we have $(a, b') = (p \wedge y, p \vee y)$ and $a \sqsubseteq y \sqsubseteq b'$ (by Lemma 35). Condition (15b) for pairs $a \sqsubseteq p$ and $a \sqsubseteq y$ gives $a \sqsubseteq p \vee y = b'$, and so p, y are \diamond -neighbors.
- $x \rightarrow y \leftarrow x'$. By Lemma 29, nodes x, x' have a common neighbor z with $d(p, z) = d(p, x) - 1 = d(p, x') - 1$. Since $z \in I(p, x)$, we must have $z \in [a, b]$ by convexity of $[a, b]$. Similarly, $z \in [a', b']$. Since (x, y, x', z) is a 4-cycle, we must have $x \leftarrow z \rightarrow x'$. Since $a = p \wedge x$ and $z \in [a, x]$, we must have $a = p \wedge z$. By a symmetric argument, $a' = p \wedge z$, and so $a = a'$. We know that $p, x, x', b, b' \in \mathcal{L}_a^+$. By Lemma 17, \mathcal{L}_a^+ is a modular semilattice which is convex in Γ . Since $y \in I(x, x')$, we must have $y = x \vee x' \in \mathcal{L}_a^+$ by convexity. Since joins $p \vee x = b, p \vee x' = b', x \vee x' = y$ exist, the join $\hat{b} = p \vee x \vee x'$ must also exist in \mathcal{L}_a^+ by definition of modular semilattices. We obtain $a \sqsubseteq \hat{b}$ and $p, y \in [a, \hat{b}]$, and so p, y are \diamond -neighbors.

□

Lemma 53 *If p, x are \diamond -neighbors then $(x, p \diamond q, q)$ is a shortest subpath.*

Proof We use induction on $d(x, p \diamond q)$. Let x' be a median of $x, p \diamond q, q$. First, assume that $x' \neq x$. By Lemma 52, p, x' are \diamond -neighbors (since $x' \in I(x, p \diamond q)$). By the induction hypothesis, $(x', p \diamond q, q)$ is a shortest subpath. (x, x', q) is also a shortest subpath, and thus so is $(x, x', p \diamond q, q)$. Now assume that $x' = x$, i.e. $(p \diamond q, x, q)$ is a shortest subpath. $(p, p \diamond q, q)$ is a shortest subpath, and thus so is $(p, p \diamond q, x, q)$. Let $(a, b) = (p \wedge x, p \vee x)$. (p, a, x) and (p, b, x) are shortest subpaths, and thus so are (p, a, x, q) and (p, b, x, q) , implying $a, b \in I(p, q)$. (a, x, b) is a shortest subpath since $a \leq x \leq b$, and so x is a median of a, b, q . Thus, we must have $x = a \vee_{pq} b$. We also have $a \sqsubseteq b$, and thus $a \in [p \nabla q, p]$ by Eq. (15). By a similar argument, $b \in [p, p \triangle q]$. We have $a \leq_{pq} p \nabla q$ and $b \leq_{pq} p \triangle q$, and so $x = a \vee_{pq} b \leq_{pq} (p \nabla q) \vee_{pq} (p \triangle q) = p \diamond q$. This implies that $(x, p \diamond q, q)$ is a shortest subpath. □

Lemma 54 *If (p, u, q) is a shortest subpath then the following sequences are also shortest subpaths:*

- (i) $(p, p \triangle q, u \triangle q, q)$;
- (ii) $(p, p \nabla q, u \nabla q, q)$;
- (iii) $(p, p \diamond q, u \diamond q, q)$.

Proof First, let us show (i). For brevity, denote $b = p \triangle q \in I(p, q)$. Define $z = b \vee_{pq} u$. (b, z, u) is a shortest subpath (since z is the median of b, u, q). Also, $u \preceq_{pq} z$, meaning that (p, u, z, q) is a shortest subpath. Lemma 35 for elements p, b, u, z gives that $u \sqsubseteq z$. This shows that $z \in \mathcal{L}_u^+ \cap I(u, q)$, and thus $(u, z, u \triangle q, q)$ is a shortest subpath by Eq. (15). Since $(z, u \triangle q, q)$ and (p, b, z, q) are shortest subpaths, so is $(p, b, z, u \triangle q, q)$.

A symmetric argument gives (ii). We can now show (iii) as follows: $u \diamond q = (u \nabla q) \vee_{uq} (u \triangle q) = \text{median}(u \nabla q, u \triangle q, q) = (u \nabla q) \vee_{pq} (u \triangle q) \succeq_{pq} (p \nabla q) \vee_{pq} (p \triangle q) = p \diamond q$. \square

Theorem 55 *If (p_0, p_1, \dots, p_k) is a \diamond -path with $p_0 = p$ then $(p_i, p \diamond^i q, q)$ is a shortest subpath for any $i \in [0, k]$.*

Proof We use induction on i . For $i = 0$ the claim holds by Lemma 53. Let us show it for $i \in [k]$. Denote $\bar{p} = p_{i-1}$, $x = p_i$, $u = p \diamond^{i-1} q$. By the induction hypothesis, $(p_{i-1}, p \diamond^{i-1} q, q) = (\bar{p}, u, q)$ is a shortest subpath. By Lemma 53, $(x, \bar{p} \diamond q, q)$ is a shortest subpath. By Lemma 54, $(\bar{p} \diamond q, u \diamond q, q)$ is a shortest subpath. Thus, $(x, \bar{p} \diamond q, u \diamond q, q)$ and so $(x, u \diamond q, q) = (p_i, p \diamond^i q, q)$ are shortest subpaths. \square

Corollary 56 *The normal p - q path $(p, p \diamond^1 q, p \diamond^2 q, \dots, q)$ is a shortest \diamond -path from p to q .*

Proof Let (p_0, p_1, \dots, p_k) be a shortest \diamond -path with $(p_0, p_k) = (p, q)$. By Theorem 55 with $q = p_k$, $(p_k, p \diamond^k q, q) = (q, p \diamond^k q, q)$ is a shortest subpath, therefore $p \diamond^k q = q$. This means that the length of the normal p - q path is at most k . \square

Our next goal is to prove that set $B_k^\diamond(p)$ is convex in Γ for any p and k . For that we will first need a technical result.

Lemma 57 *Suppose that (p, p', q', q) is an isometric rectangle in an extended modular complex Γ and p, q are \diamond -neighbors. Then p', q' are also \diamond -neighbors.*

Proof Note that $I(p, q) \cup I(p', q') \subseteq I(p, q') \cap I(p', q)$. For a node $x \in I(p, q)$ let $x' \in I(p', q')$ be the median of x, p', q' . (This median is unique and equals $x \vee_{pq'} p'$, since $x, p' \in I(p, q')$. Furthermore, p' is a median of p, p', q' and q' is a median of q, p', q' , so this notation is consistent). Since (p, x, q, q') and (x, x', q') are shortest subpaths, so is (p, x, x', q') . Analogously, since (p', p, x, q) and (p', x', x) are shortest subpaths, so is (p', x', x, q) . We can now conclude that (p, p', x', x) and (q, q', x', x) are isometric rectangles.

Denote $(a, b) = (p \wedge q, p \vee q)$, then $a \sqsubseteq b$ and $a, b \in I(p, q)$. Since (p, p', a', a) is an isometric rectangle and $a \sqsubseteq p$, we get $a' \sqsubseteq p'$ by Lemma 35. By similar arguments we get $p' \sqsupseteq a' \sqsubseteq q'$ and $p' \sqsubseteq b' \sqsupseteq q'$. (p', a', q') and (p', b', q') are shortest subpaths, and thus $(a', b') = (p' \wedge q', p' \vee q')$. We have $a' \sqsubseteq b'$ by condition (15b), and so p', q' are \diamond -neighbors. \square

Theorem 58 For each $p \in \Gamma$ and $k \geq 0$, set $B_k^\diamond(p)$ is convex in Γ .

Proof We use induction on k . Consider $k \geq 1$. By Lemma 29, it suffices to show that for distinct nodes $q_1, q_2 \in B_k^\diamond(p)$ and a common neighbor q' of q_1, q_2 we have $q' \in B_k^\diamond(p)$. Denote $p_1 = p \diamond^{k-1} q_1$ and $p_2 = p \diamond^{k-1} q_2$. By Theorem 55, (p_1, p_2, q_2) and (p_2, p_1, q_1) are shortest subpaths. Let p' be a median of p_1, p_2, q' . Since $p' \in I(p_1, p_2)$, we have $d^\diamond(p, p') \leq k - 1$ by the induction hypothesis. It thus suffices to show that p', q' are \diamond -neighbors. We can assume that $p_1 \neq p_2$ (otherwise $p_1 = p_2 = p'$ and the claim holds by Lemma 52). By symmetry, we can assume that $p_1 \neq p'$. We know that (p_2, p', p_1, q_1) and (p_1, p', q') are shortest subpaths. We also have $D \stackrel{\text{def}}{=} d(p_1, p') \geq 1$, $d(p', q') + D \leq d(p_1, q_1) + 1$ and $d(p_1, q_1) + D \leq d(p', q') + 1$. This implies that $D = 1$, $d(p_1, q_1) = d(p', q')$ and (p_1, p', q', q_1) is an isometric rectangle. Lemma 57 now gives that p', q' are \diamond -neighbors, as desired. \square

To conclude this section, we state one property of operations $\triangle, \nabla, \diamond$ that will be needed later.

Lemma 59 Consider elements p, u, q and operation $\circ \in \{\triangle, \nabla, \diamond\}$ such that $(p, p \circ q, u, q)$ is a shortest subpath. Then $p \circ u = p \circ q$.

Proof To see the claim for $\circ = \triangle$, observe that for every $x \in \mathcal{L}_p^+$ sequence $(p, x, p \triangle q, u, q)$ is a shortest subpath and thus there exists a shortest x - u path going through $p \triangle q$. This means that $p \triangle q$ is the gate of u at \mathcal{L}_p^+ , i.e. $p \triangle u = p \triangle q$. A symmetric argument gives the claim for $\circ = \nabla$. Now suppose that $(p, p \diamond q, u, q)$ is a shortest subpath. We know that $(p, p \nabla q, p \diamond q)$ and thus $(p, p \nabla q, p \diamond q, u, q)$ are shortest subpaths, so the previous result gives $a \stackrel{\text{def}}{=} p \nabla u = p \nabla q$. Similarly, we have $b \stackrel{\text{def}}{=} p \triangle u = p \triangle q$. By definition, $m = p \diamond u$ is the unique median of a, b, u . It can be seen that m is also a median of a, b, q . (Note that (a, u, q) , (a, m, u) and thus (a, m, u, q) are shortest subpaths). Therefore, $m = p \diamond q$. \square

Remark 4 For nodes p, q with $d^\diamond(p, q) = k \geq 1$ let us define $p \blacklozenge q = p \diamond^{k-1} q$. It can be deduced from Theorem 55 that $p \blacklozenge q$ is the gate of q at the convex set $B_{k-1}^\diamond(p)$ where $k = d^\diamond(p, q)$. This operation was introduced in [5] under the name “ Δ -gate of p at q ” (in the case when $\sqsubseteq = \stackrel{\text{BP}}{\sqsubseteq}$). The normal p - q path (p_0, p_1, \dots, p_k) was also used in [5] (and was shown to be a shortest \diamond -path from p to q). However, it was defined by a different construction, namely via a recursion $p_{i-1} = p \blacklozenge p_i$ for $i = k, k-1, \dots, 1$. To our knowledge, operations $\nabla, \triangle, \diamond$ and Theorem 55 have not appeared in [5]; instead, the proof techniques in [5] relied on the notion of “Helly graphs” and on the operation $\langle\langle x, y \rangle\rangle$ (the minimal convex set containing x, y), which we do not use.

5.7 Connectivity of $\text{dom } f$

In this section we will show that following.

Theorem 60 If $f : \Gamma \rightarrow \overline{\mathbb{R}}$ is an L -convex function on an extended modular complex Γ and $p, q \in \text{dom } f$ then $p \triangle q, p \nabla q \in \text{dom } f$.

Corollary 61 For any $p, q \in \Gamma$ there exists a p - q path (p_0, p_1, \dots, p_k) in Γ^\square with $(p_0, p_k) = (p, q)$ which is a shortest subpath and has the following property: for any L -convex function f on Γ with $p, q \in \text{dom } f$ one has $p_i \in \text{dom } f$ for all $i \in [0, k]$. (This path is obtained by setting either $p_{i+1} = p_i \nabla q$ or $p_{i+1} = p_i \triangle q$; one of these two elements is distinct from p_i if $p_i \neq q$).

To prove Theorem 60, we start with technical observations.

Lemma 62 Consider elements $p, q, u \in \text{dom } f$.

- (a) If $p \sqsupseteq u \sqsubseteq q$ then $p \triangle q \in \text{dom } f$. Similarly, if $p \sqsubseteq u \sqsupseteq q$ and $p, q \in \text{dom } f$ then $p \nabla q \in \text{dom } f$.
- (b) If $p \sqsubseteq u \sqsubseteq q$ then $p \triangle q \in \text{dom } f$. Similarly, if $p \sqsupseteq u \sqsupseteq q$ then $p \nabla q \in \text{dom } f$.

Proof By symmetry, it suffices to prove only the first claim in (a) and in (b). In both cases we denote $a = p \triangle q$.

(a) Denote $s = p \wedge q$, then $p \sqsupseteq s \sqsubseteq q$. We know that f is submodular on \mathcal{L}_s^+ . It suffices to show that $a \in \mathcal{E}(p, q)$; this would imply that $a \in \text{dom } f$.

Point p has coordinates $v_{pq}(p) = (X, 0)$ where $X = \mu(s, p)$. We have $a \in I(p, q)$, $p \sqsubseteq a$ and $a \wedge p = p$, therefore a has coordinates $v_{pq}(a) = (X, \dots)$. Suppose that $a \notin \mathcal{E}(p, q)$, then there exists $a' \in I(p, q)$ with $v_{pq}(a') = (X, \dots)$ and $\mu(p, a') > \mu(p, a)$. We have $a' \wedge p \in [s, p]$ and $\mu(s, a' \wedge p) = X = \mu(s, p)$, thus $a' \wedge p = p$ or equivalently $p \leq a'$. For any $x \in I(p, q)$ we have $s \sqsubseteq x$ (this follows from condition $p \sqsupseteq s \sqsubseteq q$, Lemma 24(b) and properties (15a,b)). Thus, we have $s \sqsubseteq a'$ and hence $p \sqsubseteq a'$ (since $s \leq p \leq a'$). Equation (15) gives $a' \in [p, p \triangle q] = [p, a]$, contradicting condition $\mu(p, a') > \mu(p, a)$.

(b) We know that f^* is submodular on \mathcal{L}_u^* . Note that $[x, y] \in I([p, u], [u, q])$ if and only if $x \sqsubseteq y$, $x \in [p, u]$, $y \in [u, q]$ (by Lemma 37). For such $[x, y]$ we have $v_{[p, u], [u, q]}([x, y]) = (\mu(x, u), \mu(u, y))$. It suffices to show that $[p, a] \in \mathcal{E}([p, u], [u, q])$; this would imply that $[p, a] \in \text{dom } f^*$ and $a \in \text{dom } f$.

From the lemma's assumption and Eq. (15) we conclude that $p \sqsubseteq u \sqsubseteq a \sqsubseteq q$ and $p \sqsubseteq a$, and hence $[p, a] \in \mathcal{L}_u^* \cap I([p, u], [u, q])$. Elements $[p, u]$ and $[p, a]$ have coordinates respectively $(X, 0)$ and (X, Y) where $X = \mu(p, u)$ and $Y = \mu(u, a)$. If $[p, a] \notin \mathcal{E}([p, u], [u, q])$ then there exists element $[p, a'] \in \mathcal{L}_u^* \cap I([p, u], [u, q])$ whose second coordinate is $Y' = \mu(u, a') > \mu(u, a) = Y$. This contradicts Eq. (15). \square

Lemma 63 Consider elements $p, q, u \in \Gamma$ such that $u \in [p, p \triangle q]$. Define $u^+ = u \triangle q$ and $u^- = u \nabla q$. Then $p \triangle q = p \triangle u^+$ and $p \nabla q = p \nabla u^-$.

Proof Clearly, we have $u \sqsubseteq p \triangle q$, and sequence $(p, u, p \triangle q, q)$ is a shortest subpath. Thus, we have $p \triangle q \in \mathcal{L}_u^+ \cap I(u, q)$, and so $p \triangle q \in [u, u \triangle q] = [u, u^+]$ by Eq. (15). Sequence $(p, u, p \triangle q, u^+, q)$ is thus a shortest subpath, so Lemma 59 gives the first claim: $p \triangle q = p \triangle u^+$.

To show the second claim, let $x = u \vee_{pq} (p \nabla q)$ be the median of $u, p \nabla q, q$. We have $x \sqsubseteq u$ (by the same argument as in the proof of Lemma 51) and $x \in I(u, q)$,

therefore $x \in [u \vee q, u] = [u^-, u]$ by Eq. (15). Sequence (u, x, u^-, q) is a thus shortest subpath. $(p, p \vee q, x, q)$ is also a shortest subpath, and hence so is $(p, p \vee q, x, u^-, q)$. Lemma 59 now gives $p \vee q = p \vee u^-$. \square

Lemma 64 *If $p, q \in \text{dom } f$ then $p \triangle q, p \vee q \in \text{dom } f$.*

Proof We use induction on $d(p, q)$. Suppose that $d(p, q) > 0$. By Lemma 49, there exists $u \in I(p, q) \cap \text{dom } f$ such that $pu \in \Gamma^\square$. By symmetry, we can assume that $p \sqsubset u$. We then have $u \in [p, p \triangle q] - \{p\}$. Denote $u^+ = u \triangle q$ and $u^- = u \vee q$, then we have $u^+, u^- \in \text{dom } f$ by the induction hypothesis.

Let us show that $p \vee q \in \text{dom } f$. We have $p \vee q = p \vee u^-$ by Lemma 63. If $u^- \neq q$ then the claim holds by the induction hypothesis. If $u^- = q$ then $u \sqsupseteq q$, and the claim follows from Lemma 62(a).

Let us show that $p \triangle q \in \text{dom } f$. We have $p \triangle q = p \triangle u^+$ by Lemma 63. If $u^+ \neq q$ then the claim holds by the induction hypothesis. If $u^+ = q$ then $u \sqsubseteq q$, and the claim follows from Lemma 62(b). \square

5.8 Proof of Theorem 20

Theorem 20 (restated). *Consider extended modular complexes Γ, Γ' and functions $f, f' : \Gamma \rightarrow \overline{\mathbb{R}}$ and $\tilde{f} : \Gamma \times \Gamma' \rightarrow \overline{\mathbb{R}}$.*

- (a) *If f, f' are L -convex on Γ then $f + f'$ and $c \cdot f$ for $c \in \mathbb{R}_+$ are also L -convex on Γ .*
- (b) *If f is L -convex on Γ and $\tilde{f}(p, p') = f(p)$ for $(p, p') \in \Gamma \times \Gamma'$ then \tilde{f} is L -convex on $\Gamma \times \Gamma'$.*
- (c) *If \tilde{f} is L -convex on $\Gamma \times \Gamma'$ and $f(p) = \tilde{f}(p, p')$ for fixed $p' \in \Gamma'$ then f is L -convex on Γ .*
- (d) *The indicator function $\delta_U : V \rightarrow \{0, \infty\}$ is L -convex on Γ in the following cases:*
 - (i) *U is a d_Γ -convex set;*
 - (ii) *$U = \{p, q\}$ for elements p, q with $p \sqsubseteq q$.*
- (e) *Function $\mu_\Gamma : V \times V \rightarrow \mathbb{R}_+$ is L -convex on $\Gamma \times \Gamma$.*
- (f) *If f is L -convex on Γ then the restriction of f to $\mathcal{L}_p^\sigma(\Gamma)$ is submodular on $\mathcal{L}_p^\sigma(\Gamma)$ for every $p \in \Gamma$ and $\sigma \in \{-, +\}$.*

(a) We need to show that $\text{dom}(f + f') = (\text{dom } f) \cap (\text{dom } f')$ is connected in Γ^\square ; checking other conditions of L -convexity is straightforward. Consider $p, q \in (\text{dom } f) \cap (\text{dom } f')$, and let (p_0, p_1, \dots, p_k) be the path in Γ^\square constructed in Corollary 61. By the corollary, for each $i \in [0, k]$ we have $p_i \in \text{dom } f$ and $p_i \in \text{dom } f'$, which implies the claim.

(b) The claim holds by Lemma 42(a).

(c) It suffices to prove that $\text{dom } f$ is connected in Γ^\square , the claim will then follow from Lemma 42(b). Consider $p, q \in \text{dom } f$. We know that $(p, p'), (q, p') \in \text{dom } \tilde{f}$, and thus there exists path $(\tilde{p}_0, \tilde{p}_1, \dots, \tilde{p}_k)$ in $(\Gamma \times \Gamma')^\square$ with $(\tilde{p}_0, \tilde{p}_k) = ((p, p'), (q, p'))$ and $\tilde{p}_i \in \text{dom } \tilde{f}$ for all i . By Corollary 61, such path can be chosen so that $\tilde{p}_i \in I((p, p'), (q, p'))$ for all i . Thus, we have $\tilde{p}_i = (p_i, p')$ for all i . Sequence (p_0, p_1, \dots, p_k) is now a path in Γ^\square satisfying $p_i \in \text{dom } f$ for all i .

(d) Clearly, in both cases U is connected in Γ^\square ; in the case (i) this holds since U is connected in Γ , and in the case (ii) the claim is trivial.

Let U^* be the set of vertices $[p, q]$ in Γ^* with $p, q \in U$, and for $x \in \Gamma$ denote $U_x^* = U^* \cap \mathcal{L}_x^*$. Note that $\delta_U^* = \delta_{U^*}$, and the restriction of δ_{U^*} to \mathcal{L}_x^* equals $\delta_{U_x^*}$. It suffices to show that U_x^* is convex in \mathcal{L}_x^* ; the submodularity of $\delta_{U_x^*}$ on \mathcal{L}_x^* will then follow from [15, Lemma 3.7(4)].

If $U = \{p, q\}$ where $p \sqsubseteq q$ then $U_x^* = \{[p, q]\}$ if $x \in [p, q]$, and $U_x^* = \emptyset$ otherwise; in both cases U_x^* is clearly convex. Suppose that U is a convex set in Γ . It suffices to show that U^* is convex in Γ^* . (Since \mathcal{L}_x^* is convex in Γ^* by Lemma 17, the intersection $U^* \cap \mathcal{L}_x^*$ would then be convex in Γ^* and thus also in \mathcal{L}_x^*).

We use Lemma 31 to show that U^* is convex in Γ^* . Let $[u, v] \in \Gamma^*$ be common neighbor of distinct $[p, q], [p', q'] \in U^*$; we need to show that $[u, v] \in U^*$. Modulo symmetry, 4 cases are possible:

- $p = u = p'$ and $q \rightarrow v \rightarrow q'$.
- $p = u = p'$ and $q \rightarrow v \leftarrow q'$.
- $p = u = p'$ and $q \leftarrow v \rightarrow q'$.
- $p \rightarrow u = p'$ and $q \rightarrow v = q'$.

In each of these cases conditions $p, q, p', q' \in U$ and convexity of U implies that $u, v \in U$, and so $[u, v] \in U^*$.

(e) The claim holds by Lemma 44(a).

(f) The claim holds by Lemma 39.

Corollary 65 *If function f is an L -convex on Γ and U is convex in Γ then function $f + \delta_U$ is L -convex on Γ . In particular, $f + \delta_{B_k^\diamond(x)}$ is L -convex for any $x \in \Gamma$ and $k \geq 0$.*

5.9 f -extremality

Given an L -convex function $f : \Gamma \rightarrow \overline{\mathbb{R}}$, we say that pair (p, q) is f -extremal if $p, q \in \text{dom } f$ and $f(p) = \min\{f(x) : x \in I(p, q)\}$. We say that it is *strictly f -extremal* if $p, q \in \text{dom } f$ and $\arg \min\{f(x) : x \in I(p, q)\} = \{p\}$. In this section we establish some key facts about f -extremal pairs that will be used in the analysis of the \diamond -SDA algorithm.

Lemma 66 *Suppose that (p, q) is f -extremal, $p \nabla q \neq p$ and $s \in \arg \min\{f(s) : s \in [p \nabla q, p] - \{p\}\}$. Then (s, q) is f -extremal.*

Proof Note that $p \nabla q \in \text{dom } f$ by Theorem 60, and thus $s \in \text{dom } f$. Suppose the claim is false, then there exists $q' \in I(s, q) - \{s\}$ with $f(q') < f(s)$. By Lemma 49, there exists $u \in I(s, q') \subseteq I(p, q)$ with $us \in \Gamma^\square$ and $f(u) < f(s)$. First, suppose that $u \sqsubset s$. Note that $u \leq s \leq p$ and $f(u) \geq f(p)$. We claim that there exists $x \in I(u, p) = [u, p]$ with $x \sqsubset p$ and $f(x) \leq f(u)$. Indeed, define sequence $u_0 \sqsubset u_1 \sqsubset \dots \sqsubset u_k \sqsubset p$ with $u_0 = u$ and $f(u_0) \geq f(u_1) \geq \dots \geq f(u_k) \geq f(p)$ via the following rule: if $u_i \sqsubset p$ then set $k = i$ and stop, otherwise let u_{i+1} be an element in $[u_i, p]$ with $u_i \sqsubset u_{i+1}$ and $f(u_i) \geq f(u_{i+1})$, which exists by Lemma 49. Taking $x = u_k$ now proves the claim.

Clearly, $x \in I(p, q)$, and thus $x \in [p \nabla q, p]$ and $f(x) \leq f(u) < f(s)$, which contradicts the choice of s . Thus, we must have $u \sqsupset s$. Note that $p \wedge u = s$ since $p \geq$

$s \leq u$ and (p, s, u) is a shortest subpath. We know that f is submodular on a modular semilattice \mathcal{L}_s^+ , and $p, u \in \mathcal{L}_s^+$. Thus, $f(p) + f(u) \geq f(s) + \sum_{a \in \mathcal{E}(p, u)} \lambda_a f(a)$ where λ_a are nonnegative numbers with $\sum_{a \in \mathcal{E}(p, u)} \lambda_a = 1$. This gives a contradiction since $f(u) < f(s)$ and $f(p) \leq f(a)$ for all $a \in \mathcal{E}(p, u) \subseteq I(p, q)$. \square

Corollary 67 *If (p, q) is f -extremal and $p \wedge q$ exists then $(p \wedge q, q)$ is f -extremal.*

Proof The claim follows by repeatedly applying Lemma 66 (with an induction argument). \square

Lemma 68 *If (p, q) is f -extremal then $f(p \diamond q) \leq f(q)$.*

Proof We use induction on $d(p, q)$. We can assume that $p \diamond q \neq q$, otherwise the claim is trivial. It would suffice to show the following:

(\star) *There exists $y \neq q$ such that $f(y) \leq f(q)$ and $(p \diamond q, y, q)$ is a shortest subpath.*

Indeed, by Lemma 59 we would then have $p \diamond q = p \diamond y$, and so the induction hypothesis for (p, y) would give $f(p \diamond q) = f(p \diamond y) \leq f(y) \leq f(q)$.

One of $p \vee q, p \triangle q$ must be different from p . By symmetry, we can assume that $p \vee q \neq p$. We can further assume the following: if $p \vee q \neq p$ and $p \triangle q \neq p$ then condition $p \sqsupseteq p \wedge q \sqsubseteq q$ does **not** hold, where the existence of $p \wedge q$ is a part of the condition. Indeed, conditions $p \sqsupseteq p \wedge q \sqsubseteq q$ and $p \sqsubseteq p \vee q \sqsupseteq q$ cannot hold simultaneously (otherwise we would have $p \diamond q = q$), so we can pick the condition that does not hold and then change the orientation if necessary.

Denote $x = p \vee q$, and pick $s \in \arg \min\{f(s) : s \in [x, p] - \{p\}\}$. By Lemma 66, (s, q) is f -extremal. Four cases are possible.

Case 1: $d^\diamond(s, q) \geq 2$ Denote $y = s \diamond q$, then $y \neq q$. By Lemma 54, $(p \diamond q, y, q)$ is a shortest subpath. By the induction hypothesis for (s, q) we have $f(y) \leq f(q)$. Thus, (\star) holds.

In the remaining cases we have $d^\diamond(s, q) \leq 1$. We will denote $a = s \wedge q = p \wedge q$. Note that $p \sqsupseteq s \sqsupseteq a \sqsubseteq q$. Furthermore, $x \in [a, s]$. Indeed, (p, s, x, q) and thus (s, x, q) are shortest subpaths and $s \sqsupseteq x$, so $x \geq s \vee q = s \wedge q = a$ by Eq. (15). One can check that (p, s, x, a, q) is a shortest subpath (since (p, s, q) , (s, a, q) , (s, x, a) are shortest subpaths). Furthermore, $(p \wedge q, q) = (a, q)$ is f -extremal by Corollary 67.

Case 2: $d^\diamond(s, q) \leq 1, p \sqsupseteq a$ We have $p \sqsupseteq a \sqsubseteq q$, which implies that $p \triangle q = p$ and thus $p \diamond q = p \vee q = x$. f -extremality of (a, q) gives $f(a) \leq f(q)$. We have $a \neq q$ since we assumed that $p \diamond q \neq q$. Thus, element $y = a$ satisfies condition (\star).

Case 3: $d^\diamond(s, q) \leq 1, p \not\sqsupseteq a < q$ We have $p \diamond a = p \vee a = p \vee q = x$ where the first two equalities hold since $p \geq a$ and by Lemma 59, respectively. Since (a, q) is f -extremal, we get $a \in \text{dom } f$. Since $a \neq q$, we can apply the induction hypothesis for (p, a) and get $f(x) = f(p \diamond a) \leq f(a)$.

Conditions $d^\diamond(s, q) \leq 1$ and $x \in [a, s]$ imply that $y \stackrel{\text{def}}{=} x \vee q$ exists and $a \sqsubseteq x$. Note that $x \wedge q = a$. By submodularity inequality for x, q in modular semilattice \mathcal{L}_a^+ we get $f(x) + f(q) \geq f(a) + f(y)$, and so $f(y) \leq f(q)$. Clearly, (p, x, y, q)

is a shortest subpath. By applying Lemma 24 to modular semilattice \mathcal{L}_a^\uparrow we obtain that $(p, p \triangle q, y, q)$ is a shortest subpath (since $y \in I(p, q)$ and $y \geq q$). Since $y \geq_{pq} x = p \nabla q$ and $y \geq_{pq} p \triangle q$, we get $y \geq_{pq} (p \nabla q) \vee_{pq} (p \triangle q) = p \diamond q$, i.e. $(p, p \diamond q, y, q)$ is a shortest subpath. We have $x \neq a$ (since $p \supseteq x$ and $p \not\supseteq a$) and thus $y \neq q$. Therefore, (\star) holds.

Case 4: $d^\diamond(s, q) \leq 1$, $p \not\supseteq a = q$ We have $q \leq x \leq s \leq p, q \sqsubseteq s$ (since $d^\diamond(s, q) \leq 1$) and $x \sqsubseteq p$. Note that $p \diamond q = p \nabla q = x$. Also, $q \neq x$ (since $q \not\sqsubseteq p$ and $x \sqsubseteq p$). By Theorem 60 we have $x \in \text{dom } f$. We claim that (\star) holds. Indeed, suppose this is false, then $f(y) > f(q)$ for all $y \in [q, x] - \{q\}$. Denote $\alpha = [q, s]$, $\beta = [x, s]$, $\gamma = [x, p]$, then $\alpha, \beta, \gamma \in \mathcal{L}_s^*$ and $\alpha \wedge \gamma = \alpha \cap \gamma = \beta$. We know that function $f^* : \mathcal{L}_s^* \rightarrow \mathbb{R}$ is submodular on \mathcal{L}_s^* . Condition $f(q) < f(x)$ implies that $f^*(\alpha) < f^*(\beta)$. Condition $f(p) \leq f(s)$ implies that $f^*(\gamma) \leq f^*(\beta)$.

Consider $\delta = [u, v] \in \mathcal{E}(\alpha, \gamma) - \{\gamma\}$. By Lemma 37 we must have $u \in [q, x]$ and $v \in [s, p]$. We must have $v \neq p$ (otherwise condition $\delta \neq \gamma$ would imply that $u \in [q, x] - \{x\}$; we would also have $u \sqsubseteq v = p$ which is impossible since $x = p \nabla q$). Conditions $u \in [q, x]$ and $v \in [s, p] - \{p\}$ have two implications: (i) the second coordinate of $v_{\alpha\gamma}(\delta)$ is strictly smaller than the second coordinate of $v_{\alpha\gamma}(\gamma)$; (ii) $f(u) \geq f(q)$, $f(v) \geq f(s)$ and hence $f^*(\delta) \geq f^*(\alpha)$.

By submodularity, $f^*(\alpha) + f^*(\gamma) \geq f^*(\beta) + \sum_{\delta \in \mathcal{E}(\alpha, \gamma)} \lambda_\delta f^*(\delta)$ where λ_δ are nonnegative numbers with $\sum_{\delta \in \mathcal{E}(\alpha, \gamma)} \lambda_\delta = 1$. By implication (i), we have $\lambda_\gamma > 0$. The submodularity inequality can be rewritten as

$$[f^*(\gamma) - f^*(\beta)] + \sum_{\delta \in \mathcal{E}(\alpha, \gamma) : \lambda_\delta > 0} \lambda_\delta [f^*(\alpha) - f^*(\delta)] \geq 0$$

All numbers in square brackets are non-positive, thus they must all be 0. This implies that $f^*(\gamma) = f^*(\beta)$ and $f^*(\alpha) = f^*(\gamma)$, which contradicts condition $f^*(\alpha) < f^*(\beta)$. \square

We conclude this section with a couple of results that will help to deal with cases when values of f are not unique.

Lemma 69 *If (p, q) is strictly f -extremal and $p \diamond q \neq q$ then $f(p \diamond q) < f(q)$.*

Proof For $\varepsilon > 0$ define function $f_\varepsilon : \Gamma \rightarrow \overline{\mathbb{R}}$ via $f_\varepsilon(x) = f(x) + \varepsilon \mu(x, q)$. Note that f_ε is L -convex on Γ . Clearly, there exists $\varepsilon > 0$ such that (p, q) is f_ε -extremal. Using Lemma 68, we obtain $f(p \diamond q) = f_\varepsilon(p \diamond q) - \varepsilon \mu(p \diamond q, q) < f_\varepsilon(p \diamond q) \leq f_\varepsilon(q) = f(q)$. \square

Lemma 70 *Suppose that (p, q) is f -extremal, $p \leq q$, $p \not\sqsubseteq q$, $q \diamond p \leq p \diamond q$ and $f(q') \geq f(q)$ for all $q' \in [p \diamond q, q]$. Then there exists $p' \in [p, q \diamond p] - \{p\}$ with $f(p') = f(p)$.*

Proof Denote $(a, b) = (q \diamond p, p \diamond q)$, then $p \leq a \leq b \leq q$, $p \sqsubseteq b$, $a \sqsubseteq q$ and hence $a \sqsubseteq b$. By Theorem 60, we have $a, b \in \text{dom } f$. Denote $\alpha = [p, b]$, $\beta = [a, b]$, $\gamma = [a, q]$, then $\alpha, \beta, \gamma \in \mathcal{L}_a^*$ and $\alpha \wedge \gamma = \beta$. Note that $f^*(\alpha) + f^*(\gamma) = f^*(\beta) +$

$f(p) + f(q)$. By submodularity, $f^*(\alpha) + f^*(\gamma) \geq f^*(\beta) + \sum_{\delta \in \mathcal{E}(\alpha, \gamma)} \lambda_\delta f^*(\delta)$, or equivalently $f(p) + f(q) \geq \sum_{\delta \in \mathcal{E}(\alpha, \gamma)} \lambda_\delta f^*(\delta)$ where λ_δ are nonnegative numbers with $\sum_{\delta \in \mathcal{E}(\alpha, \gamma)} \lambda_\delta = 1$. For every $\delta = [u, v] \in \mathcal{E}(\alpha, \gamma)$ we have $u \in [p, a]$, $v \in [b, q]$ by Lemma 37. Furthermore, $f(u) \geq f(p)$ and $f(v) \geq f(q)$ by the lemma's assumption, and so $f^*(\delta) \geq f(p) + f(q)$. This implies that $f(u) = f(p)$, $f(v) = f(q)$ for every $\delta = [u, v] \in \mathcal{E}(\alpha, \gamma)$ with $\lambda_\delta > 0$. Note that $[p, q] \notin \mathcal{L}_a^*$ since $p \not\sqsubseteq q$, and so $\alpha \vee \gamma$ does not exist in \mathcal{L}_a^* . This implies that $\alpha \neq \beta \neq \gamma$. Therefore, there exists $\delta = [u, v] \in \mathcal{E}(\alpha, \gamma) - \{\alpha\}$ with $\lambda_\delta > 0$ (and hence with $f(u) = f(p)$). We claim that $u \neq p$ (and so $u \in [p, a] - \{p\}$, implying the lemma). Indeed, if $u = p$ then condition $\delta \neq \alpha$ gives that $v \neq b$, and hence $v \in [b, q] - \{b\}$. This is impossible since $p = u \sqsubseteq v$ and $b = p \diamond q = p \triangle q$. \square

5.10 Proof of Theorem 22: analysis of the \diamond -SDA algorithm

Theorem 22 (restated). *Let Γ be an extended modular complex and $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ be an L -convex function on Γ^n . \diamond -SDA algorithm applied to function f terminates after generating exactly $1 + \max_{i \in [n]} d_\Gamma^\diamond(x_i, \text{opt}_i(f))$ distinct points, where x is the initial vertex and $\text{opt}_i(f)$ is as defined in Theorem 12.*

First, we show the following fact.

Lemma 71 *Suppose that Γ is a Cartesian product of extended modular complexes: $\Gamma = \Gamma_1 \times \dots \times \Gamma_n$. Then $d_\Gamma^\diamond(x, y) = \max_{i \in [n]} d_{\Gamma_i}^\diamond(x_i, y_i)$.*

Proof It suffices to show the following fact for extended modular complexes Γ, Γ' .

(*) *The following conditions are equivalent for elements $x, y \in \Gamma$ and $x', y' \in \Gamma'$:*

- (a) $(x, x'), (y, y')$ are \diamond -neighbors;
- (b) x, y are \diamond -neighbors and x', y' are \diamond -neighbors.

Let us define $(a, b) = (x \wedge y, x \vee y)$ and $(a', b') = (x' \wedge y', x' \vee y')$. (These expressions are defined if either (a) or (b) holds). By the definition of the Cartesian product $\Gamma \times \Gamma'$, we have the following implications:

- (a) $\Leftrightarrow (x, x') \sqsupseteq (a, a') \sqsubseteq (y, y')$ and $(x, x') \sqsubseteq (b, b') \sqsupseteq (y, y')$
- $\Leftrightarrow x \sqsupseteq a \sqsubseteq y$ and $x' \sqsupseteq a' \sqsubseteq y'$ and $x \sqsubseteq b \sqsupseteq y$ and $x' \sqsubseteq b' \sqsupseteq y' \Leftrightarrow$ (b)

\square

In the light of this lemma, it suffices to prove Theorem 22 in the case when $n = 1$. Accordingly, from on we assume that \diamond -SDA is applied to minimize function $f : \Gamma \rightarrow \overline{\mathbb{R}}$ which is L -convex on an extended modular complex Γ .

Let \bar{x} be the initial vertex of the \diamond -SDA algorithm, and for $k \geq 0$ define function $f_k = f + \delta_{B_k^\diamond(\bar{x})}$. By Corollary 65, f_k is L -convex on Γ for any $k \geq 0$.

Lemma 72 *For any $q \in \arg \min f_{k-1}$ there exists $p \in \arg \min f_k$ with $d^\diamond(p, q) \leq 1$.*

Proof Let p be an element of $\arg \min f_k$ which is closest to q . Clearly, (p, q) is f_k -extremal, and in fact strictly f_k -extremal. (If there exists $p' \in I(p, q) - \{p\}$ with $f_k(p') = f_k(p)$ then p' also belongs to $\arg \min f_k$ and is closer to q than p , contradicting the choice of p). Since $d^\diamond(\bar{x}, p) \leq k$, there exists $u \in B_{k-1}^\diamond(\bar{x})$ with $d^\diamond(u, p) \leq 1$. $(u, p \diamond q, q)$ is a shortest subpath by Theorem 55. We have $u, q \in B_{k-1}^\diamond(\bar{x})$ and thus $p \diamond q \in B_{k-1}^\diamond(\bar{x})$ by convexity of $B_{k-1}^\diamond(\bar{x})$ (Theorem 58). If $p \diamond q \neq q$ then $f_k(p \diamond q) < f_k(q)$ by Lemma 69, which contradicts the choice of q . Thus, we must have $p \diamond q = q$ and so $d^\diamond(p, q) \leq 1$. \square

Lemma 73 Consider $x \in \text{dom } f$, and let p be an element of $\arg \min \{f(y) \mid y \in B_1^\diamond(x)\}$ which is closest to x . Let $x^- \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^-(\Gamma)\}$ and $x^+ \in \arg \min \{f(y) \mid y \in \mathcal{L}_x^+(\Gamma)\}$. Then $x^- \sqsubseteq p \sqsubseteq x^+$.

Proof By symmetry, it suffices to prove that $p \sqsubseteq q$ where we denoted $q = x^+$. Define $(a, b) = (p \wedge x, p \vee x)$. Note that $[a, b]$ is a modular lattice. Since $q \geq x$, the meet $s = p \wedge q$ exists and satisfies $s \in [a, p]$. Define $y = s \vee x$, then $y \in [s, q]$ and $p \vee y = b$. Since $p \sqsupseteq a \sqsubseteq x$, we get $p \sqsupseteq s \sqsubseteq y$ (using, in particular, Lemma 35). This means that $p \nabla q = s$, $p \triangle q \geq b$ and $p \diamond q \in [y, q] \subseteq [x, q]$.

Note that $p, q \in B_1^\diamond(x)$ and thus $I(p, q) \subseteq B_1^\diamond(x)$ by convexity of $B_1^\diamond(x)$. The choice of p thus implies that (p, q) is f -extremal. We claim that $f(p) = f(s)$. Indeed, suppose that $f(p) < f(s)$. Pair (s, q) is f -extremal by Corollary 67, and thus $f(p) < f(s) \leq f(q)$. Let p' be an element of $I(p, q)$ with $f(p') = f(p)$ which is closest element to q . By construction, $p' \neq q$ and (p', q) is strictly f -extremal. If $p' \diamond q \neq q$ then $f(p' \diamond q) < f(q)$ by Lemma 69, and $(p, p \diamond q, p' \diamond q, q)$ is a shortest subpath by Lemma 54. Condition $p \diamond q \in [x, q]$ thus implies that $x \leq p \diamond q \leq p' \diamond q \leq q$. Condition $x \sqsubseteq q$ now gives $x \sqsubseteq p' \diamond q$, contradicting the choice of q . Thus, we must have $p' \diamond q = q$. By the choice of p' we have $f(p') < f(p' \wedge q)$, so the submodularity inequality for (p', q) gives that $f(p' \vee q) < f(q)$ where $p' \vee q$ exists since $p' \diamond q = q$. Note that $x \leq q \leq p' \vee q$ and $p' \vee q \in I(p, q) \subseteq B_1^\diamond(x)$. This implies that $x \sqsubseteq p' \vee q$, contradicting the choice of q . We proved that $f(p) = f(s)$, and thus $p = s$ by the choice of p , since (p, s, x) is a shortest subpath.

We now know that $p \sqsubseteq y \sqsubseteq q$ where $y = b = x \vee p$. Suppose that $p \not\sqsubseteq q$. Clearly, we have $p \diamond q \in [y, q]$, $q \diamond p \in [p, y]$ and $q \diamond p \leq y \leq p \diamond q$. By choice of q we must have $f(q') \geq f(q)$ for any $q' \in [p \diamond q, q] \subseteq [y, q] \subseteq [x, q]$. By Lemma 70, there exists $p' \in [p, q \diamond p] - \{p\} \subseteq [p, y] - \{p\}$ with $f(p') = f(p)$. This contradicts the choice of p , since (p, p', y, x) is a shortest subpath. \square

Corollary 74 Element x^\diamond computed in line 4 of Algorithm 2 satisfies $x^\diamond \in \arg \min \{f(y) \mid y \in B_1^\diamond(x)\}$.

Proof Clearly, any $y \in \mathcal{L}_{x^-}^+(\Gamma) \cap \mathcal{L}_{x^+}^-(\Gamma)$ satisfies $y \in B_1^\diamond(x)$. The claim now follows directly from Lemma 73. \square

Lemma 72 and Corollary 74 imply that after k iterations point x in \diamond -SDA satisfies $x \in \arg \min f_k$ (via an induction argument). This yields Theorem 22.

6 VCSP proofs

In this section we prove two results: Theorem 23 (BLP relaxation solves language Φ_Γ for an extended modular complex Γ) and the hardness direction of Theorem 3. Below we give some background on Valued Constraint Satisfaction Problems (VCSPs) which will be needed for these proofs.

Let us fix finite set D . Let $\mathcal{O}^{(m)}$ be the set of operations $g : D^m \rightarrow D$. Operation $g \in \mathcal{O}^{(m)}$ is called *symmetric* if $g(x_1, \dots, x_m) = g(x_{\pi(1)}, \dots, x_{\pi(m)})$ for any tuple $(x_1, \dots, x_m) \in D^m$ and any permutation $\pi : [m] \rightarrow [m]$. A *fractional operation of arity m* is a probability distribution over $\mathcal{O}^{(m)}$, i.e. vector $\omega \in [0, 1]^{\mathcal{O}^{(m)}}$ with $\sum_g \omega(g) = 1$. Fractional operation ω is called *symmetric* if all operations in $\text{supp}(\omega) = \{g \in \mathcal{O}^{(m)} \mid \omega(g) > 0\}$ are symmetric.

A cost function $f : D^n \rightarrow \mathbb{R}$ is said to *admit* ω (or ω is a *fractional polymorphism of f*) if

$$\sum_{g \in \text{supp}(\omega)} \omega(g) f(g(x^1, \dots, x^m)) \leq \frac{1}{m} \sum_{i=1}^m f(x^i) \quad \forall x^1, \dots, x^m \in D^n$$

where operation $g(\cdot)$ is applied componentwise, i.e.

$$g(x^1, \dots, x^m) = (g(x_1^1, \dots, x_1^m), \dots, g(x_n^1, \dots, x_n^m)) \in D^n$$

Language Φ over D is said to admit ω if all functions $f \in \Phi$ admit ω . In this case ω is called a *fractional polymorphism of Φ* .

The *expressive power* $\langle \Phi \rangle$ of language Φ is defined as the set of all cost functions $f : D^n \rightarrow \mathbb{R}$ of the form $f(x) = \min_{y \in D^k} f_{\mathcal{I}}(x, y)$ where \mathcal{I} is a Φ -instance with $n + k$ variables. It is known that if Φ admits a fractional polymorphism ω then $\langle \Phi \rangle$ also admits ω .

Definition 75 Language Φ on domain D is said to satisfy condition (MC) if there exist distinct $a, b \in D$ and binary function $f \in \langle \Phi \rangle$ such that $\arg \min f = \{(a, b), (b, a)\}$.

Theorem 76 ([34, Theorem 3.4], [25, Theorem 5]) *Let Φ be a finite-valued language on domain D such that for every $a \in D$ there exists a unary cost function $g_a \in \langle \Phi \rangle$ with $\arg \min g_a = \{a\}$. If Φ does not satisfy (MC) then Φ admits a symmetric fractional polymorphism of every arity $m \geq 2$.*

Theorem 77 ([25, Theorem 1, Proposition 8]) *Let Φ be a general-valued language. BLP solves Φ if and only if Φ admits a symmetric fractional polymorphism of every arity $m \geq 2$. If this condition holds then an optimal solution of any Φ -instance can be computed in polynomial time.*

6.1 Proof of Theorem 23

Recall that Φ_Γ is the language over domain $D = V_\Gamma$ that consists of all functions $f : D^n \rightarrow \mathbb{R}$ such that f is L -convex on Γ^n .

Theorem 78 *If Γ is an extended modular complex then Φ_Γ does not satisfy (MC).*

Proof Suppose the claim is false, then there exists an instance $f : \Gamma \times \Gamma \times \Gamma^n \rightarrow \overline{\mathbb{R}}$ of Φ_Γ with $n + 2 \geq 2$ variables such that function $g : \Gamma \times \Gamma \rightarrow \overline{\mathbb{R}}$ defined via $g(x, y) = \min_{z \in \Gamma^n} f(x, y, z)$ satisfies $\arg \min g = \{(a, b), (b, a)\}$ for some distinct $a, b \in \Gamma$. Denote $A = \arg \min f \subseteq \Gamma \times \Gamma \times \Gamma^n$, $A_{ab} = \{(a, b, z) \mid z \in \Gamma^n\}$ and $A_{ba} = \{(b, a, z) \mid z \in \Gamma^n\}$. By construction, $A \subseteq A_{ab} \cup A_{ba}$, $A \cap A_{ab} \neq \emptyset$ and $A \cap A_{ba} \neq \emptyset$. By Lemma 48, the subgraph of $(\Gamma \times \Gamma \times \Gamma^n)^\square$ induced by set A is connected. Thus, there must exist $p \in A_{ab}, q \in A_{ba}$ such that $pq \in (\Gamma \times \Gamma \times \Gamma^n)^\square$. By symmetry, we can assume that $p \sqsubset q$. We have $p = (a, b, x)$ and $q = (b, a, y)$ for some $x, y \in \Gamma^n$. Condition $p \sqsubset q$ thus implies that $a \sqsubset b$ and $b \sqsubset a$, which is impossible. \square

We can finally prove that BLP solves Φ_Γ .

Theorem 79 *Let Γ be an extended modular complex. Language Φ_Γ admits a symmetric fractional polymorphism of every arity $m \geq 2$. Consequently, BLP relaxation solves Φ_Γ , and an optimal solution of any Φ_Γ -instance can be computed in polynomial time.*

Proof If language Φ_Γ were finite-valued then the claim would immediately follow from Theorems 78 and 76. The main concern will thus be dealing with non finite-valued languages.

Let us define $\Phi_\Gamma^\circ = \{f \in \langle \Phi_\Gamma \rangle \mid f \text{ is finite-valued}\}$. Since language Φ_Γ does not satisfy (MC), languages $\langle \Phi_\Gamma \rangle$ and Φ_Γ° also do not satisfy (MC). Clearly, for each $a \in \Gamma$ language Φ_Γ° contains unary function g_a with $\arg \min g_a = \{a\}$, namely $g_a(x) = \mu(a, x)$. By Theorem 76, Φ_Γ° admits a symmetric fractional polymorphism of every arity $m \geq 2$. Let ω_m be such fractional polymorphism. We claim that Φ_Γ also admits ω_m . Indeed, consider function $f : \Gamma^n \rightarrow \overline{\mathbb{R}}$ in Φ_Γ . We can assume w.l.o.g. that $f(x) \geq 0$ for all $x \in \text{dom } f$, and $\mu(x, y) \geq 1$ for distinct $x, y \in \Gamma$. Define function $\mu_n : \Gamma^n \times \Gamma^n \rightarrow \mathbb{R}$ via $\mu_n(x, y) = \mu_{\Gamma^n}(x, y)$. By Theorem 44, μ_n is L -convex on $\Gamma^n \times \Gamma^n$, and thus $\mu_n \in \Phi_\Gamma$. For value $C > 0$ define function $f_C : \Gamma^n \rightarrow \mathbb{R}$ via $f_C(x) = \min_{y \in \Gamma^n} (f(y) + C\mu_n(x, y))$, then $f_C \in \Phi_\Gamma^\circ$ and thus f_C admits ω_m . Clearly, for any $C > \max_{x \in \text{dom } f} f(x)$ the following holds: $f_C(x) = f(x)$ if $x \in \text{dom } f$, and $f_C(x) \geq C$ if $f(x) = \infty$. By taking the limit $C \rightarrow \infty$ we conclude that f also admits ω_m . \square

6.2 Proof of the hardness direction of Theorem 3

We will use a technique from [26].

Consider language Φ on domain D . A pair of elements $(a, b) \in D \times D$ is called *conservative* if there exists a unary function $g_{ab} \in \langle \Phi \rangle$ with $\arg \min g_{ab} = \{a, b\}$. Let $\mathcal{S}(\Phi) \subseteq D \times D$ be the set of conservative pairs in Φ . For a tuple $p = (a, b) \in \mathcal{S}(\Phi)$, we denote $\bar{p} = (b, a)$; clearly, $\bar{p} \in \mathcal{S}(\Phi)$. Now consider two tuples $p = (a, b)$ and $q = (c, d)$. We say that pair (p, q) is *strictly submodular* if there exists binary cost function $f \in \langle \Phi \rangle$ such that

$$f(a, c) + f(b, d) < f(a, d) + f(b, c)$$

Clearly, if (p, q) is strictly submodular, then (q, p) is also strictly submodular, since function f' defined via $f'(x, y) = f(y, x)$ also belongs to $\langle \Phi \rangle$. We thus say that $\{p, q\}$ is strictly submodular if (p, q) is strictly submodular (or equivalently, if (q, p) is strictly submodular). Let $\mathcal{E}(\Phi) \subseteq \binom{\mathcal{S}(\Phi)}{2}$ be the set of strictly submodular pairs $\{p, q\}$, and define undirected graph $\mathcal{G}(\Phi) = (\mathcal{S}(\Phi), \mathcal{E}(\Phi))$.

Theorem 80 [26, Theorem 3.2(a) and Lemma 5.1(b)] *Suppose that language Φ is finite-valued and that, for each $a \in D$, there exists unary cost function $g_a \in \langle \Phi \rangle$ with $\arg \min g_a = \{a\}$.*

- (a) *If $\{p, q\}, \{q, r\} \in \mathcal{E}(\Phi)$, then $\{p, r\} \in \mathcal{E}(\Phi)$.*
- (b) *If $\{p, \bar{p}\} \in \mathcal{E}(\Phi)$ for some $p \in \mathcal{S}(\Phi)$, then Φ is NP-hard.*

Remark 5 In [26] Kolmogorov and Živný formulated Theorem 80 only in the case when Φ contains all possible $\{0, 1\}$ -valued unary functions (and thus $\mathcal{S}(\Phi) = D \times D$). However, the proofs of the results above use only weaker preconditions stated in Theorem 80.

Note that the graph defined in [26] had an edge $\{p, q\}$ if and only if our graph $\mathcal{G}(\Phi)$ has an edge $\{p, \bar{q}\}$. We translated the results from [26] accordingly.

We now apply these results to the generalized minimum 0-extension problem for metric space (V, μ) and subset $F \subseteq \binom{V}{2}$. Let us define the following language over domain $D = V$:

$$\Phi = \{\mu\} \cup \{\mu_a : a \in V\} \cup \{\mu_{ab} : \{a, b\} \in F\}$$

where unary function μ_a, μ_{ab} are defined via $\mu_a(x) = \mu(x, a)$ and $\mu_{ab}(x) = \min\{\mu(x, a), \mu(x, b)\}$. Clearly, we have $\Phi \subseteq \langle 0\text{-Ext}[\mu, F] \rangle$.

Let $H = H_\mu = (V, E, w)$ be the graph corresponding to μ . Suppose that H is not F -orientable; our goal is to show that Φ is NP-hard. We can assume w.l.o.g. that H is modular, otherwise Φ is NP-hard by Theorem 1. Define

$$\vec{E} = \{(a, b) : \{a, b\} \in E\} \quad \vec{F} = \{(a, b) : \{a, b\} \in F\}$$

Let us introduce relations $\parallel, \triangleleft, \approx$ for tuples $p = (a, b)$ and $q = (c, d)$ as follows:

- $p \parallel q$ if $p, q \in \vec{E}$ and (a, b, d, c) is a 4-cycle in H ;
- $p \triangleleft q$ if $p \in \vec{E}, q \in \vec{F}, p \neq q$, and (c, a, b, d) is a shortest subpath in H ;
- $p \approx q$ if at least one of the following holds: (i) $p \parallel q$; (ii) $p \triangleleft q$; (iii) $q \triangleleft p$.

Note that relation \approx is symmetric, and accordingly, $(\vec{E} \cup \vec{F}, \approx)$ is an undirected graph.

Lemma 81 *Let $p, q \in \mathcal{S}(\Phi)$.*

- (a) $\vec{E} \cup \vec{F} \subseteq \mathcal{S}(\Phi)$.
- (b) *If $p \parallel q$, then $\{p, q\} \in \mathcal{E}(\Phi)$.*
- (c) *If $p \triangleleft q$, then $\{p, q\} \in \mathcal{E}(\Phi)$.*
- (d) *If $p \approx q$, then $\{p, q\} \in \mathcal{E}(\Phi)$.*

Proof (a) First, $\vec{F} \subseteq \mathcal{S}(\Phi)$ holds by the definition of Φ .

As for \vec{E} , for each $\{a, b\} \in E$, we construct a function $g_{ab} \in \langle \Phi \rangle$ by $g_{ab}(x) := \mu_a(x) + \mu_b(x)$. From the minimality of H , we have $I(a, b) = \{a, b\}$. We therefore see that $\arg \min g_{ab} = \{a, b\}$, and thus, we obtain $(a, b) \in \mathcal{S}(\Phi)$.

(b) Suppose that $p = (a, b) \in \vec{E}$, $q = (c, d) \in \vec{E}$, and (a, b, d, c) is a 4-cycle in H . The modularity of H implies that $\{a, d\} \notin E$; otherwise a, b, d would not have a median. Similarly, $\{b, c\} \notin E$; otherwise b, d, c would not have a median. By Theorem 6(a), sequences (a, b, d) and (b, d, c) are shortest subpaths in H . This implies that $\mu(a, c) + \mu(b, d) < \mu(a, d) + \mu(b, c)$. We thereby obtain $\{p, q\} \in \mathcal{E}(\Phi)$.

(c) Suppose that $p = (a, b) \in \vec{E}$, $q = (c, d) \in \vec{F}$, and (c, a, b, d) is a shortest subpath in H . Then $\mu(a, c) + \mu(b, d) < (\mu(a, c) + \mu(a, b)) + (\mu(a, b) + \mu(b, d)) = \mu(b, c) + \mu(a, d)$, and so, we again obtain $\{p, q\} \in \mathcal{E}(\Phi)$.

(d) We decompose $p \approx q$ into cases. If $p \parallel q$, we apply (b). If $p \triangleleft q$, we apply (c). If $q \triangleleft p$, we utilize $\{p, q\} = \{q, p\}$ and apply (c) again. We conclude that $\{p, q\} \in \mathcal{E}(\Phi)$. \square

We can finally finish the proof of Theorem 3.

We claim that there exists an element $p \in \vec{E} \cup \vec{F}$ such that p, \bar{p} are connected by a path in $(\vec{E} \cup \vec{F}, \approx)$. Indeed, if no such p exists, then there exists a mapping $\sigma : \vec{E} \cup \vec{F} \rightarrow \{-1, +1\}$ with the following properties: (i) if $q \approx r$, then $\sigma(q) = \sigma(r)$; (ii) $\sigma(q) = -\sigma(\bar{q})$ for each $q \in \vec{E} \cup \vec{F}$. Such mapping can be constructed by greedily assigning connected components of graph $(\vec{E} \cup \vec{F}, \approx)$; the assumption ensures that no conflicts arise. Using the mapping σ , we can define an orientation of (H, F) by orienting $\{a, b\} \in E \cup F$ as $a \rightarrow b$ if $\sigma(a, b) = +1$ and as $a \leftarrow b$ if $\sigma(a, b) = -1$. Clearly, this orientation is admissible, which contradicts the assumption that H is not F -orientable.

From Theorem 80(a) and Lemma 81, we conclude that $\{p, \bar{p}\} \in \mathcal{E}(\Phi)$. Theorem 80(b) now gives that Φ is NP-hard.

7 Open questions

In this section we state some open problems that we find interesting.

One of the motivations for this work was trying to find the “structure” of tractable finite-valued languages. This leads to the following question.

- Can it be case that for every tractable finite-valued language Φ there exists an extended modular complex Γ such all functions $f \in \Phi$ are L -convex on Γ ? A negative answer could yield new interesting classes of functions that are not yet known in the area of Discrete Convex Analysis.

To simplify the problem, one may consider a restricted class of languages.

- Resolve the question above for (a) languages of the form $0\text{-Ext}[\mu, F]$ where μ is a metric on V and F is a subset of 2^V ; (b) languages corresponding to directed metrics, as in [17]. Note that [17] characterized tractable directed metrics on a star, with some complicated fractional polymorphisms. As a special case, it would be interesting to verify whether these tractable classes are captured by the theory developed in this paper. (This question was suggested by an anonymous reviewer).

There are also some open problems concerning L -convex functions on extended modular complexes.

- Is it possible to show that the normal SDA algorithm terminates after a polynomial number of iterations?
- As shown in [15, 16], if f an an L -convex function on a modular complex Γ then f^* is L -convex on Γ^* . Is this also true for extended modular complexes?
- Let $\sqsubseteq, \sqsubseteq'$ be admissible relations for modular complex Γ such that \sqsubseteq coarsens \sqsubseteq' . Is it true that if function f is L -convex on (Γ, \sqsubseteq) then it is also L -convex on (Γ, \sqsubseteq') ? This is known to hold when Γ is an oriented path and $(\sqsubseteq, \sqsubseteq') = (\sqsubseteq, \stackrel{\text{bp}}{\sqsubseteq})$: any $L^{\mathbb{Q}}$ -convex function is submodular on an integer lattice. However, we were unable to prove or disprove it in the general case.
- Theorem 23 shows that L -convex functions on extended modular complexes admit a binary symmetric fractional polymorphism, however the proof of this fact is non-constructive. Is there a more explicit construction of this polymorphism? As remarked in Section 6 of [15], answering this question might require a further thorough investigation of orientable modular graphs.

Acknowledgements We thank the anonymous reviewers for their careful reading of our manuscript and their many insightful comments and suggestions.

Funding Open access funding provided by Institute of Science and Technology (IST Austria).

Declarations

Conflict of interest The authors do not have any conflicting interests related to this work.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Bandelt, H.-J.: Networks with condorcet solutions. *Eur. J. Oper. Res.* **20**, 314–326 (1985)
2. Bandelt, H.-J., van de Vel, M., Verheul, E.: Modular interval spaces. *Math. Nachr.* **163**, 177–201 (1993)
3. Birkhoff, G.: *Lattice Theory*, 3rd edn. American Mathematical Society, (1967)
4. Blake, A., Kohli, P., Rother, C.: *Markov Random Fields for Vision and Image Processing*, MIT Press, (2011)
5. Chalopin, J., Chepoi, V., Hirai, H., Osajda, D.: Weakly modular graphs and nonpositive curvature. *Memoirs AMS*, 268(1309), (2020)
6. Chekuri, C., Khanna, S., Naor, J., Zosin, L.: A linear programming formulation and approximation algorithms for the metric labeling problem. *SIAM J. Dis. Math.* **18**(3), 608–625 (2004)
7. Chepoi, V.: Classification of graphs by means of metric triangles (in Russian). *Metody Diskretnogo Analiza* **49**, 75–93 (1989)
8. Chuzhoy, J., Naor, J.: The hardness of metric labeling. *SIAM J. Comput.* **36**(2), 498–515 (2006)
9. Dress, A.W.M., Scharlau, R.: Gated sets in metric spaces. *Aequationes Math.* **34**, 112–120 (1987)

10. Dvorak, M., Kolmogorov, V.: Generalized minimum 0-extension problem and discrete convexity. [arXiv:2109.10203v3](https://arxiv.org/abs/2109.10203v3), October (2021)
11. Felzenszwalb, P., Pap, G., Tardos, E., Zabih, R.: Globally optimal pixel labeling algorithms for tree metrics. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), (2010)
12. Fujishige, S.: Submodular Funct. Optimiz., 2nd edn. Elsevier, Amsterdam (2005)
13. Fujishige, S., Murota, K.: Notes on L-/M-convex functions and the separation theorems. *Math. Progr. Ser. A* **88**, 129–146 (2000)
14. Gridchyn, I., Kolmogorov, V.: Potts model, parametric maxflow and k-submodular functions. In: Proceedings of the 14th IEEE International Conference on Computer Vision (ICCV'13), pp. 2320–2327. IEEE, (2013)
15. Hirai, H.: Discrete convexity and polynomial solvability in minimum 0-extension problems. *Math. Progr. Ser. A* **155**, 1–55 (2016)
16. Hirai, H.: L-convexity on graph structures. *J. Oper. Res. Soc. Japan* **61**, 71–109 (2018)
17. Hirai, H., Mizutani, R.: Minimum 0-extension problems on directed metrics. *Discret. Optim.* **40**, (2021)
18. Huber, A., Kolmogorov, V.: Towards Minimizing k -Submodular Functions. In: Proceedings of the 2nd International Symposium on Combinatorial Optimization (ISCO'12), volume 7422 of *Lecture Notes in Computer Science*, pp. 451–462. Springer, (2012)
19. Huber, A., Krokhin, A., Powell, R.: Skew bisubmodularity and valued CSPs. *SIAM J. Comput.* **43**(3), 1064–1084 (2014)
20. Karzanov, A.V.: Minimum 0-extensions of graph metrics. *Eur. J. Comb.* **19**, 71–101 (1998)
21. Karzanov, A.V.: One more well-solved case of the multifacility location problem. *Discret. Optim.* **1**(1), 51–66 (2004)
22. Kleinberg, J., Tardos, É.: Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *J. ACM* **49**, 616–639 (2002)
23. Kolmogorov, V.: Submodularity on a tree: Unifying l^\sharp -convex and bisubmodular functions. In: Proceedings of the 36th International Symposium on Mathematical Foundations of Computer Science (MFCS'11), volume 6907 of *Lecture Notes in Computer Science*, pp. 400–411. Springer, (2011)
24. Kolmogorov, V., Shioura, A.: New algorithms for convex cost tension problem with application to computer vision. *Discret. Optim.* **6**, 378–393 (2009)
25. Kolmogorov, V., Thapper, J., Živný, S.: The power of linear programming for general-valued CSPs. *SIAM J. Comput.* **44**(1), 1–36 (2015)
26. Kolmogorov, V., Živný, S.: The complexity of conservative valued CSPs. *J. ACM*, **60**(2), Article 10 (2013)
27. Murota, K.: Discrete convex analysis. *Math. Progr.* **83**, 313–371 (1998)
28. Murota, K.: *Discret. Convex Anal.* SIAM, Philadelphia (2003)
29. Murota, K.: On steepest descent algorithms for discrete convex functions. *SIAM J. Optim.* **14**, 699–707 (2003)
30. Murota, K., Shioura, A.: M-convex function on generalized polymatroid. *Math. Oper. Res.* **24**, 95–105 (1999)
31. Murota, K., Shioura, A.: Exact bounds for steepest descent algorithms of L-convex function minimization. *Oper. Res. Lett.* **42**, 361–366 (2014)
32. Murota, K., Tamura, A.: Proximity theorems of discrete convex functions. *Math. Progr. Ser. A* **99**, 539–562 (2004)
33. Tansel, B.C., Francis, R.L., Lowe, T.J.: Location on networks I II. *Manag. Sci.* **29**, 498–511 (1983)
34. Thapper, J., Živný, S.: The complexity of finite-valued CSPs. *J. ACM (JACM)*, **63**(4), (2016)
35. van de Vel, M.L.J.: *Theory Convex Struct.* North-Holland, Amsterdam (1993)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.