Pictures to

"Metrics with finite sets of primitive extensions"

Earlier necessary and sufficient conditions have been found for the path metrics d^H that admit no primitive extensions except themselves.

Theorem 1.1 [10]. *Let H be a connected graph. Then $|\Pi(\mathrm{d}^H)| = 1$ if and only if H is bipartite, orientable and contains no isometric k-cycle with $k \geq 6$.*

Here a *k-cycle* is a (simple) circuit C_k on k nodes (considered as a closed path or as a graph depending on the context). A graph H is *orientable* if the edges of H can be oriented so that for any 4-cycle $C = (v_0, e_1, v_1, ..., e_4, v_4 = v_0)$, the orientations of the opposite edges e_1 and e_3 are different along the cycle, and similarly for e_2 and e_4 (a feasible orientation is depicted in Fig. 2). A subgraph H' of H is *isometric* if $\mathrm{d}^{H'}(xy) = \mathrm{d}^H(xy)$ for any nodes x, y of H'. The graphs H as in Theorem 1.1 are called *frames*. We refer to [10] for other results on frames, in particular, those related to a generalization of the multiterminal cut problem from [6].

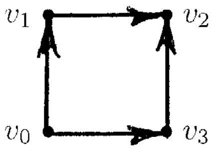

Fig. 1

CIn this paper we focus on the case when $\Pi(\mu)$ is finite. Such a metric μ is called *primitively finite*, or a *PF-metric*. (Our initial motivation for studying this case came up from the multiflow (multicommodity flow) area. Extreme extensions arise as optimal dual solutions in one sort of multiflow problems where one is asked for maximizing the sum of flow values weighted by a given metric μ on the set of terminals. The finiteness of $\Pi(\mu)$ means that, up to the similarity, the number of unavoidable optimal dual solutions occurring in the problem instances concerning this μ is finite. These aspects will be discussed in Section 6.) Our main result is the following.

Theorem 1.2. *Let μ be a positive rational metric on a finite set T. The following are equivalent:*

(i) *$\Pi(\mu)$ is finite;*

(ii) *there exist a frame H and an integer $\lambda > 0$ such that $\lambda\mu$ is a submetric of d^H;*

(iii) *the least generating graph $G = (V, E)$ for a modular closure (V, m) of μ is bipartite and contains as an isometric subgraph neither C_k with $k \geq 6$ nor $K_{3,3}^-$.*

We have to explain the notions used in (iii) of this theorem that gives a combinatorial characterization of the PF-metrics.

First, for a metric m on V, a point $v \in V$ is called a *median* of a triple $\{s_0, s_1, s_2\}$ in V if

$$m(s_i v) + m(v s_j) = m(s_i s_j) \quad \text{for all } 0 \leq i < j \leq 2. \tag{1.1}$$

When a median exists for each triple, m is called *modular*. By a *modular closure* of μ we mean a certain its extension (V,m) which is modular and is constructed by the following process. Initially set $V := T$ and $m := \mu$. Choose in V a triple $\{s_0, s_1, s_2\}$ without a median, add a new point v to V and define the distances from v to the s_i's so as to satisfy (1.1) (such distances exist and are unique). Then define distances from v to the other points in V as follows. Let $V' \subset V$ be the set of points of which distances from v have already been defined; initially $V' = \{s_0, s_1, s_2, v\}$. Choose a point $u \in V - V'$ and put

$$m(uv) := \max\{m(ux) - m(xv) \,:\, x \in V' - \{v\}\}. \tag{1.2}$$

Update $V' := V' \cup \{u\}$ and iterate until $V' = V$. One can see that m remains an extension of μ. Repeat this procedure for a next medianless triple in the current (V,m), and so on.

When the process terminates, the resulting (V,m) has medians for all triples and is just the desired extension of μ. Note that a priori m depends on the order in which the medianless triples are treated (however, it is invariant when μ is primitively finite, as we explain in Section 5). We show in Section 2 that for a rational metric μ the above process does terminate in a finite number of steps and, moreover, that the resulting m is a primitive extension of μ.

Second, a spanning subgraph G of K_V is said to *generate* m if m coincides with the path metric defined by G whose each edge e has length $m(e)$. The *least graph generating*, or, briefly, *LG-graph*, of m is obtained by deleting all redundant edges from K_V, where xy is *redundant* if there is $z \in V - \{x,y\}$ with $m(xz) + m(zy) = m(xy)$ (such a z is said to be *between* x and y).

Third, $K^-_{3,3}$ is the graph obtained by deleting one edge from $K_{3,3}$, where $K_{p,q}$ is the complete bipartite graph whose parts (i.e., the maximal stable sets) consist of p and q nodes; see Fig. 2. We call a bipartite graph without isometric subgraphs C_k, $k \geq 6$, and $K^-_{3,3}$ a *semiframe*. Note that $K^-_{3,3}$ is non-orientable, so every frame is a semiframe. On the other hand, every graph $K_{p,q}$ is a semiframe, but it is not a frame if $p, q \geq 3$. In a bipartite graph an induced subgraph C_6 or $K^-_{3,3}$ is, obviously, isometric.

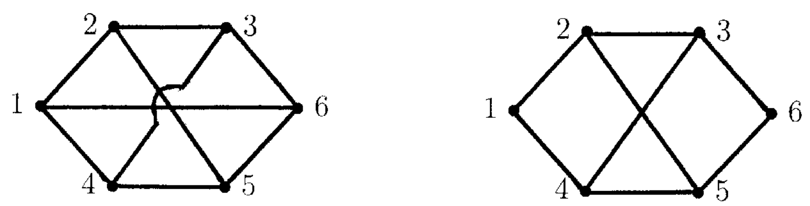

Fig. 2 (a) $K_{3,3}$ (b) $K^-_{3,3}$

We now briefly outline the method of proof of Theorem 1.2. Implication (ii)→(i) will follow from Theorem 1.1 and a rather simple fact that if μ is a submetric of a to multiflows.

Among a variety of tools used in our proofs, we, in particular, apply results of Bandelt on hereditary modular graphs. A graph H is called *modular* if d^H is modular, and *hereditary modular* if each isometric subgraph of H is modular. In particular, any modular graph is bipartite.

Theorem 1.8 [3]. *Let $H = (T, U)$ be a graph.*

(i) *H is hereditary modular if and only if H is bipartite and contains no isometric k-cycle with $k \geq 6$.*

(ii) *If H is modular but not hereditary modular, then H contains an isometric 6-cycle that, in its turn, is contained in a (not necessarily induced) cube in H (see Fig. 3).*

(iii) *If H is bipartite but not modular, then H contains a medianless triple $\{s_0, s_1, s_2\}$ with $\mathrm{d}^H(s_0s_1) = \mathrm{d}^H(s_0s_2) \geq 2$ and $\mathrm{d}^H(s_1s_2) = 2$.*

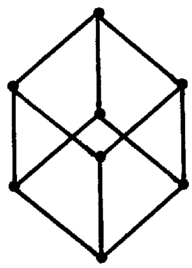

Fig. 3

In view of (i), the frames (semiframes) are exactly those hereditary modular graphs which are orientable (respectively, without induced subgraphs $K_{3,3}^-$).

A majority of results in this paper can be extended, with a due care, to arbitrary, not necessarily rational and finite, metrics μ; we, however, consider only the finite rational case to make our description shorter and technically simpler.

2. The modular closure and least generating graph

Let m be an extension of μ to V. A sequence $P = (x_0, x_1, \ldots, x_k)$ of points of V is called a *path on* V; we refer to P as a *T-path* if $x_0, x_k \in T$, and a *cycle* if $x_0 = x_k$. For brevity we write $P = x_0x_1 \ldots x_k$. The *length* of P with respect to m, or the *m-length*, is $m(P) = m(x_0x_1) + \ldots + m(x_{k-1}x_k)$, and P is called *m-shortest* if $m(P) = m(x_0x_k)$. The set of m-shortest T-paths is denoted by $\mathcal{G}(m) = \mathcal{G}(T, m)$.

If extensions m, m', m'' of μ to V satisfy $m \geq \lambda m' + (1 - \lambda)m''$ with $0 < \lambda \leq 1$, we say that *m' decomposes m*. So m is extreme if and only if no $m' \neq m$ decomposes m.

relation is symmetric and transitive, and we call a maximal set of dependent edges an *orbit* of G. Then (3.1) implies that

(3.2) the distances $m(e)$ of all edges e of an orbit of G are the same.

Now suppose that G contains an induced subgraph $H' = (T', U')$ isomorphic to $K_{3,3}^-$ (notation $H' \simeq K_{3,3}^-$). Note that H' is isometric since G is bipartite. Moreover, it is easy to check that all edges of H' are dependent (via 4-cycles in H'), therefore, by (3.2),

(3.3) the submetric μ' of m to T' is $\lambda d^{H'}$ for some $\lambda > 0$.

A feature of $K_{3,3}^-$ is that its metric $d = d^{K_{3,3}^-}$ has a primitive extension which, in its turn, has a proper submetric d' isomorphic to $\frac{1}{2}d$. Then d' can be extended in a similar way, and one can repeat such a procedure as many times as one wishes, every time obtaining a new primitive extension of the initial metric, due to (i) and (ii) in Statement 3.1. More precisely, the following is true (cf. Theorem 1.3).

Statement 3.2 [10]. *For $H' = (T', U') \simeq K_{3,3}^-$, there exists a bipartite graph $G'' = (V'', E'')$ with $V'' \supset T'$ such that: (i) $m'' = \frac{1}{2}d^{G''}$ is a primitive extension of $d^{H'}$, and (ii) G'' contains $K_{3,3}^-$ as an induced subgraph.*

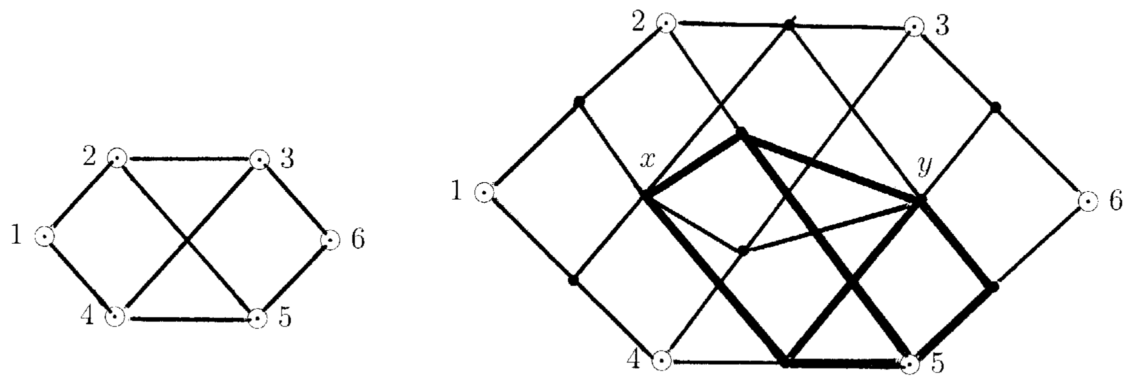

Fig. 4 (a) $H' \simeq K_{3,3}^-$ (b) G''

The desired graph G'' is drawn in Fig. 4b where for convenience the nodes of H' are labelled by $1, \ldots, 6$ as indicated in Fig. 4a. This G'' is obtained by splitting each edge $e = ij$ of H' into two edges iz_e and z_ej in series, and adding: (a) two extra nodes x and y, (b) edges xz_e for all $e = ij \in U'$ with $i, j \leq 5$, and (c) edges yz_e for all $e = ij \in U'$ with $i, j \geq 2$. An induced subgraph $K_{3,3}^-$ in G'' is drawn bold in Fig. 4b. It is not difficult to check that m'' is an extension of $d^{H'}$ and, moreover, that m'' is a tight extension of $d^{H'}$ (e.g., x and y belong to a shortest path of length six in G'' which connects nodes 1 and 4). However, a direct verification of the primitivity of m''

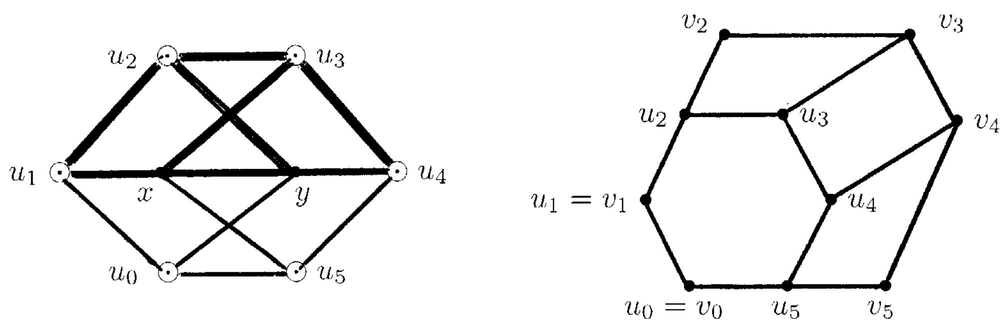

Fig. 5 (a) Γ (b)

Now assume that the α, β, γ are not the same, $\alpha < \beta \leq \gamma$ say. Let ρ be the distance function for a set W of ten points $u_0 = v_0, u_1 = v_1, v_2, \ldots, v_5, u_2, \ldots, u_5$, defined by

$$
\begin{aligned}
&\rho(v_i v_j) = m(v_i v_j), \\
&\rho(u_i u_j) = \alpha\varphi(i,j), \\
&\rho(u_i v_j) = m(v_0 v_j) - \alpha\varphi(0,j) + \alpha\varphi(i,j)
\end{aligned}
\tag{3.4}
$$

for $i, j = 0, \ldots, 5$, where $\varphi(i', j') = \mathrm{d}^C(v_{i'} v_{j'})$. One can check that ρ is indeed a metric on W (Fig. 5b illustrates the LG-graph for ρ). Moreover, by (3.4), for each $i = 0, \ldots, 5$, the ρ-length of the path $P_i = v_i u_i u_{i+1} u_{i+2} u_{i+3} v_{i+3}$ is equal to

$$
\begin{aligned}
&\rho(v_i u_i) + \rho(u_i u_{i+3}) + \rho(u_{i+3} v_{i+3}) \\
&\qquad = (m(v_0 v_i) - \alpha\varphi(0,i)) + 3\alpha + (m(v_0 v_{i+3}) - \alpha\varphi(0, i+3)) \\
&\qquad\qquad = m(v_i v_{i+3}) + 3\alpha - \alpha\varphi(i, i+3) = m(v_i v_{i+3}) \quad (= \alpha + \beta + \gamma),
\end{aligned}
$$

taking indices modulo 6 and using the facts that $m(v_0 v_i) + m(v_0 v_{i+3}) = m(v_i v_{i+3})$ and $\varphi(0,i) + \varphi(0, i+3) = 3$. So P_i is a ρ-shortest $\widetilde{T}$-path. Since each two points of W occur in some P_i, ρ is a tight extension of $\widetilde{\mu}$. But ρ is not primitive for $\widetilde{\mu}$. Nevertheless, we can use the fact that the submetric ρ' of ρ on $\{u_0, \ldots, u_5\}$ is $\alpha \mathrm{d}^{C_6}$. We further extend ρ by use of the metric $\alpha \mathrm{d}^\Gamma$ (which first extends ρ'), where Γ is the above-mentioned graph depicted in Fig. 5a. Then the resulting metric $\widetilde{\rho}$ on $W \cup \{x, y\}$ is already a primitive extension of $\widetilde{\mu}$. This follows from the observation that the collection of the above ρ-shortest paths P_i together with the shortest paths between u_i and u_{i+3} in Γ, $i = 0, 1, 2$, determines $\widetilde{\rho}$ uniquely. Again, μ has a primitive extension in which some submetric is proportional to $\mathrm{d}^{\bar{K}_{3,3}}$, providing the existence of infinitely many primitive extensions for μ. This completes the proof of implication (i)$\to$(iii) in Theorem 1.2.

4. Proof of (iii)$\to$(ii) in Theorem 1.2

Let the graph $G = (V, E)$ as above be a semiframe. We show that in this case there exists a frame H such that $\lambda\mu$ is a submetric of d^H for some $\lambda > 0$.

exists a *feasible* orientation of the edges of this orbit, i.e., each two opposite edges of any 4-cycle of G either are not oriented or have different orientations along this cycle (this matches the definition of the orientability in the Introduction). Note that Q may be orientable while the whole G is not. The graph G' has important properties exhibited in the following three lemmas.

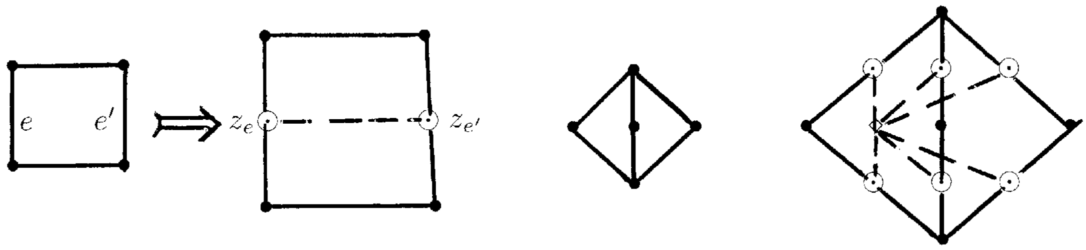

Fig. 6 (a) (b)

Lemma 4.2. *G' is a semiframe.*

Lemma 4.3. *Let $Q_1, \dots, Q_r$ be the orbits of G, and let $G' = (V', E')$ be obtained from G by the orbit splitting operation applied to Q_1. Then G' has r or $r+1$ orbits. Moreover,*

(i) for $i = 2, \dots, r$, Q_i induces an orbit Q'_i in G' formed by the edges of Q_i and all bridge-edges $z_e z_{e'}$ such that the 4-cycle in G containing e, e' has the other two edges in Q_i; also Q'_i is orientable if and only if Q_i is so;

(ii) if Q_1 is non-orientable, then Q_1 induces one orbit Q'_1 in G', which is orientable and formed by all the split- and star-edges;

(iii) if Q_1 is orientable, then Q_1 induces two orbits Q'_1 and Q''_1 in G', and the set $Q'_1 \cup Q''_1$ consists of all the split- and star-edges; also both Q'_1 and Q''_1 are orientable, and for each $e = xy \in Q_1$, one of the edges xz_e, yz_e belongs to Q'_1 while the other to Q''_1.

Lemma 4.4. *Let ρ be the length function on the edges of G defined as follows (using notation from the previous lemma):*

(4.2) (i) for $i = 2, \dots, r$, the ρ-length of each edge in Q'_i is equal to the distance $m(e)$ of an edge $e \in Q_i$;

(ii) if Q_1 is non-orientable and $e \in Q_1$, then the ρ-length of each edge in Q'_i is equal to $m(e)/2$;

(iii) if Q_1 is orientable and $e \in Q_1$, then fix an arbitrary number $0 < \alpha < m(e)$ and put $\rho(e') = \alpha$ for all $e' \in Q'_1$ and $\rho(e'') = m(e'') - \alpha$ for all $e'' \in Q''_1$.

Then the path metric $m^\rho = d^{G',\rho}$ on V' coincides with ρ on E' and is a tight extension of μ.

Γ of $\Gamma_{p,q}$ bounded by two shortest paths from $s = (0,0)$ to $t = (p,q)$ is called a *net* from s to t, or an *s–t net*. That is, Γ is a subgraph of $\Gamma_{p,q}$ induced by the nodes (i,j) satisfying $a_j \le i \le b_j$ for two sequences $0 = a_0 \le a_1 \le \ldots \le a_q \le p$ and $0 \le b_0 \le b_1 \le \ldots \le b_q = p$ with $a_j \le b_j$, $j = 0,\ldots,q$. Figure 7b illustrates a net Γ for $p = 4$ and $q = 3$. The *rightmost* (resp. *leftmost*) path from s to t in Γ is denoted by R^Γ (resp. L^Γ). Sometimes a node with coordinates (i,j) in Γ is denoted by $(i,j)_\Gamma$.

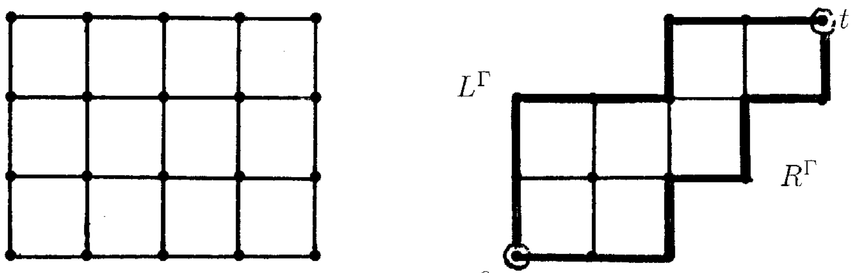

Fig. 7 (a) $\Gamma_{4,3}$ (b) Γ

We will use as an important tool the property that any two shortest paths with the same ends in G can be linked by an isometric net. More precisely,

(4.6) (i) for any $s, t \in V$ and shortest s–t paths P and P' in G, there exists an s–t net Γ in G such that $R^\Gamma = P$ and $L^\Gamma = P'$;

(ii) any 2-connected net in G is isometric.

This property was proved in [10] concerning the frames; however, the proof remains valid for the semiframes as it does not really uses the orientability of the hereditary modular graph in question but only absence of induced subgraphs $K_{3,3}^-$ in it. We outline the proof for completeness of our description.

Sketch of the proof of (4.6). We show (i) by induction on $|P|$. Let $P = sx_1\ldots x_kt$ and $P' = sy_1\ldots y_kt$. Case $|P| \le 2$ is obvious, so assume $|P| \ge 3$. Also one may assume that $x_i \ne y_j$ for $i,j = 1,\ldots,k$ (otherwise the result easily follows by induction). Since G has no isometric n-cycle with $n > 4$, there are i,j such that $d^G(x_iy_j) < i+j, 2k+2-i-j$; one may assume that $i+j \le 2k+2-i-j$, that $i \ge j$, and that $i+j$ is minimum.

Choose a shortest path $y_jz_1\ldots z_qx_i$ in G. The above assumptions and the bipartiteness of G imply that the paths $L = y_{j-1}\ldots y_1sx_1\ldots x_i$, $L' = y_{j-1}y_jz_1\ldots z_qx_i$ and $L'' = y_j\ldots y_1sx_1\ldots x_{i-1}$ are shortest (letting $x_0 = y_0 = s$). By induction there is a y_{j-1}–x_i net Γ' with $R^{\Gamma'} = L$ and $L^{\Gamma'} = L'$. Clearly L'' can be shortest only if Γ is the $(i+j-2) \times 1$ grid in which y_j, x_{i-1}, x_i, s have the coordinates $(0,1), (0,i+j-2), (1,i+j-2), (j-1,0)$, respectively (in view of $i \ge j$). Then $D = sz_{j-1}\ldots z_qx_i\ldots x_kt$ and $D' = sz_{j-1}\ldots z_1y_j\ldots y_kt$ are shortest paths in G. By induction there is a z_{j-1}–t net Γ'' with $R^{\Gamma''}$ and $L^{\Gamma''}$ to be the parts of D and D' from

from the polyhedral structure of tight spans mentioned in the Introduction).

First we recall the construction of $\mathcal{T}(\mathrm{d}^H)$ for a frame $H = (W, U)$. Each edge e of H is regarded as being homeomorphic to the closed interval (segment) $[0, 1] \subset \mathbb{R}^1$ with the natural metric σ^e on it. Each 4-cycle $C = v_0v_1v_2v_3v_0$ (considered up to reversing and cyclically shifting) is expanded into a 2-dimensional disc D^C. Formally, D^C is homeomorphic to $[0, 1] \times [0, 1] \subset \mathbb{R}^2$, the nodes v_0, v_1, v_2, v_3 are identified with the points (0,0),(0,1),(1,1),(1,0), respectively, and the edges with the corresponding segments. D^C is endowed with the ℓ_1-metric $\sigma^C = \sigma^{v_0v_1} \oplus \sigma^{v_1v_2}$, i.e., for points $x = (\xi, \eta)$ and $y = (\xi', \eta')$ in D^C, $\sigma^C(xy) = |\xi - \xi'| + |\eta - \eta'|$. If two 4-cycle $C = v_0v_1v_2v_3v_0$ and $C' = u_0u_1u_2u_3u_0$ have three common nodes, $v_i = u_i$ for $i = 0, 1, 2$ say, we identify the corresponding halves (triangles) in D^C and $D^{C'}$; namely, assuming for definiteness that v_0, v_1, v_2 are represented as (0,0),(0,1),(1,1) in both discs, respectively, we identify each point (ξ, η) for $0 \le \xi \le \eta \le 1$ in D^C with the (ξ, η) in $D^{C'}$. As a result, every bi-clique $K = (A; B)$ with $A = \{s_1, s_2\}$ and $B = \{t_1, \dots, t_k\}$ produces the shape $F(K)$, called the *folder* of K, homeomorhic to the space formed by sticking together k copies of the triangle $\{(\xi, \eta) : 0 \le \xi \le \eta \le 1\}$ along the side $\{(\alpha, \alpha) : 0 \le \alpha < 1\}$; see Fig. 8 for $k = 5$. The above metrics σ^C for 4-cycles C in K give the metric σ^K on $F(K)$. In view of Statement 4.1, two different folders have at most one vertex or one edge in common.

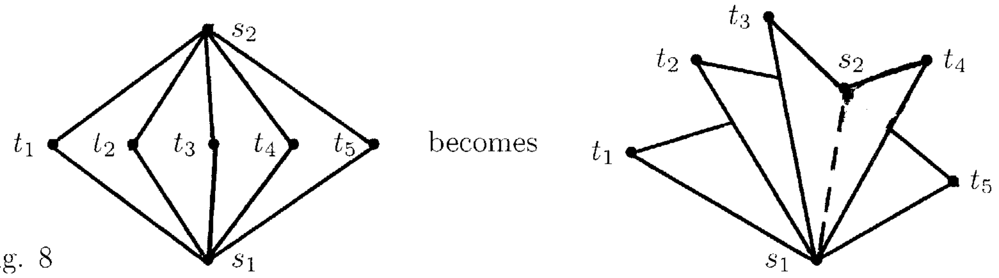

Fig. 8

The resulting space is just $\mathcal{T}(\mathrm{d}^H) = (\mathcal{X}, \sigma)$, where $\mathcal{X} = \cup(F(K) : K \in \mathcal{K}(H))$ and the global metric σ on $\mathcal{X}$ is defined in a natural way: for $x, y \in \mathcal{X}$, $\sigma(xy)$ is the infimum of values $\sigma^{q_1}(x_0x_1) + \ldots + \sigma^{q_r}(x_{r-1}x_r)$ over all finite sequences $x = x_0, x_1, \dots, x_r = y$ in which each two x_{i-1}, x_i occur in the folder $F(q_i)$ of a bi-clique q_i or in a bridge q_i.

Now suppose that h takes value 2 on the edges of some orbit Q of H, and 1 on the other edges of H. Let H' be obtained by splitting the orbit Q, taking $\alpha = 1$ (see Lemma 4.4). Since $\mathrm{d}^{H'}$ is a tight extension of $g = \mathrm{d}^{H,h}$, we have $\mathcal{T}(g) = \mathcal{T}(\mathrm{d}^{H'})$; so the folder structure of $\mathcal{T}(\mathrm{d}^{H'})$ describes $\mathcal{T}(g)$. However, we can describe $\mathcal{T}(g) = (\widehat{\mathcal{X}}, \widehat{\sigma})$ in terms of H and h themselves, as follows.

(i) Let $K = xyuvx$ be a simple bi-clique (4-cycle) in $\mathcal{K}(Q)$ with $xy, uv \in Q$ (see the definition in the beginning of Section 4). Then K induces two bi-cliques (4-cycles) $K' = xzz'vx$ and $K'' = zyuz'z$ in H', where $z = z_{xy}$ and $z' = z_{uv}$. Each of $F(K'), F(K'')$ is represented as the square $[0, 1] \times [0, 1]$, and the common segment between z and z' sticks