Morphology on Graphs and Minimum Spanning Trees

Fernand Meyer and Jean Stawiaski

MINES Paristech, CMM- Centre de morphologie mathématique, Mathématiques et Systèmes,
35 rue Saint Honoré - 77305 Fontainebleau cedex, France
{fernand.meyer,jean.stawiaksi}@mines-paristech.fr
http://cmm.ensmp.fr

Abstract. This paper revisits the construction of watershed and waterfall hierarchies through a thorough analysis of Boruvka's algorithms for constructing minimum spanning trees of edge weighted graphs. In the case where the watershed of a node weighted graph is to be constructed, we propose a distribution of weights on the edges, so that the waterfall extraction on the edge weighted graph becomes equivalent with the watershed extraction on the node weighted graph.

Keywords: Waterfall, minimum spanning tree, adjunctions on graphs.

1 Introduction

Graphs are the fundamental structures used to represent images and partitions. Different graph theoretical approaches are well suited to describe morphological segmentation: 1) detection of shortest paths or cycles for various types of distances ; 2) extraction of minimum spanning trees and forests (watershed hierarchies and segmentation with markers) ; 3) maximal flows and minimal cuts. Some links between these optimal structures have already been studied [1,2,3].

From these previous studies, it appears that minimum spanning trees and forests play a central role for watershed, waterfall and hierarchical segmentations. In this paper a thorough examination of Boruvka's algorithm, known as the first algorithm created for constructing minimum spanning trees, give us a deeper insight into the links between waterfalls and watersheds. We first define a non conventional morphology based on adjunctions between nodes and edges of weighted graphs. From these new operators, we show that waterfalls can be directly obtained from a specific morphological opening. Morphological segmentations can be well described by flooding a topographic surface. We highlight how the different steps of Boruvka's algorithm can be used to build watersheds and waterfalls on edge weighted graphs. We then address the problem of the non-uniqueness of minimal spanning trees arising when graphs have plateaux and edges having the same weight. We especially give several strategies to bypass this problem. Finally, we present an extension of this methodology in case of node weighted graphs.

M.H.F. Wilkinson and J.B.T.M. Roerdink (Eds.): ISMM 2009, LNCS 5720, pp. 161–170, 2009.

2 Adjunctions on Graphs and Interpretations in Terms of Flooding

A *non oriented graph* $G = [N, E]$ is a collection of a set N whose elements are called vertices or nodes and of a set E whose elements $u \in E$ are pairs of vertices called edges. The *weights* $[e, n]$ of the graph G are represented as two functions h and k, respectively for the edges and the nodes. h_{ij} is the weight of the edge (i, j) and k_i the weight of the node i. The same graph may have various distributions of weights. If there is no ambiguity, in the case where only one distribution of weights is considered, e will also represent the distribution of weights of the edges and n the distribution of weights of the nodes.

2.1 Various Adjunctions on Graphs

Classical morphology on graphs involves operations between nodes (i.e. operations on pixels) or operations between edges. We define in this section new types of adjunctions between both nodes and edges. These operators permit us to interpret flooding in terms of basic morphological operations.

Definition 1. *We define here several operators involving both edges and nodes of a weighted graph.*

- *an erosion* $[\varepsilon_{en} n]_{ij} = n_i \wedge n_j$ *and its adjunct dilation* $[\delta_{ne} e]_i = \bigvee_{(k \ \ neighbors\, of\, i)} e_{ik}$.
- *a dilation* $[\delta_{en} n]_{ij} = n_i \vee n_j$ *and its adjunct erosion* $[\varepsilon_{ne} e]_i = \bigwedge_{(k \ \ neighbors\, of\, i)} e_{ik}$.

Proposition 1. *The operators defined above are pairwise adjunct or dual operators:*
- ε_{ne} *and* δ_{en} *are adjunct operators.*
- ε_{en} *and* δ_{ne} *are adjunct operators.*
- ε_{ne} *and* δ_{ne} *are dual operators from the edge weight function* e *to the node weight function* n*.*
- ε_{en} *and* δ_{en} *are dual operators from the node weight function* e *to the edge weight function* n*.*

As ε_{ne} and δ_{en} are adjunct operators, the operator $\varphi_n = \varepsilon_{ne}\delta_{en}$ is a closing on n and $\gamma_e = \delta_{en}\varepsilon_{ne}$ is an opening on e.

Similarly as ε_{en} and δ_{ne} are adjunct operators, the operator $\varphi_e = \varepsilon_{en}\delta_{ne}$ is a closing on e and $\gamma_n = \delta_{ne}\varepsilon_{en}$ is an opening on n.

The operators $\varepsilon_E = \varepsilon_{en}\varepsilon_{ne}$ and $\delta_E = \delta_{en}\delta_{ne}$ are respectively the elementary erosion and the elementary dilation on a neighborhood graph operating on neighboring edges : $\overline{E} \rightarrow \overline{E}$, since $(\varepsilon_E, \delta_E)$ also form an adjunction. Likewise the operators $\varepsilon_N = \varepsilon_{ne}\varepsilon_{en}$ and $\delta_N = \delta_{ne}\delta_{en}$ are the usual elementary erosions and dilations on a neighborhood graph operating on nodes: $\overline{N} \rightarrow \overline{N}$. As such $(\varepsilon_N, \delta_N)$ also form an adjunction.

2.2 Interpretation of γ_e in Terms of Flooding

Let G represent a neighboring graph of catchment basins, such as a region adjacency graph of a watershed segmentation, then the operators defined above have a physical interpretation. $[\varepsilon nee]_i$ represents the lowest pass point between the catchment basin i and its neighbors ; its altitude or weight is the overflood level of basin i. A flooding starting in basin i would flood into the neighboring basins through this pass point. So $\gamma_e(i,j) = \delta_{en}\varepsilon_{ne}(i,j)$ represents the highest overflood level of the basins i and j. γ_e will thus be invariant on all edges which are the overflood of one of the neighboring basins or in terms of graphs the lowest edge of one of their extremities. This property is used later to show how waterfalls can be constructed locally using invariants of the opening γ_e.

3 Morphological Segmentation on Edge Weighted Graphs

Edge weighted graphs are useful in segmentation since they carry a compact representation of hierarchical segmentation : cutting all edges with a weight $> \lambda$ transforms the graph into a forest, where each tree spans a region of the segmentation. For increasing threshold values, less and less edges are cut and the associated segmentation becomes coarser. The minimum spanning tree (MST) of the edge weighted graph (among all spanning trees) is particularly useful in this context, since by cutting the edges with a weight $> \lambda$ yields the same partition as by cutting the edges above the same threshold on the initial graph.

The MST is ubiquitous in morphological segmentation issues. Segmenting with markers leads to constructing a minimum spanning forest [4], where each tree is rooted in a marker. Serge Beucher and Béatrice Marcotegui [5] have shown that the waterfall hierarchy may be obtained by constructing the watershed on the MST associated to the regional minima [6]. Jean Cousty et al [7], studying the watershed defined on edge weighted graphs, have shown the equivalence between watershed and minimum spanning forests associated to the regional minima of the graph.

3.1 Waterfalls

Waterfall hierarchies have been introduced by Serge Beucher in his thesis [8,9]. The waterfall hierarchy describes the intrication of the catchment basins and the nested structures of a topographic surface. The waterfall hierarchy may be obtained by flooding. The lowest and finest level of the hierarchy corresponds to the set of all catchment basins. The first level of hierarchy is obtained by taking the catchment basins of the topographic surface after it has been submitted to the following flooding : each catchment basin is filled up to its lowest pass point. The process is then repeated for this new topographic surface and produces the successive levels of the hierarchy. The process ends when a topographic surface is created containing only one catchment basin.

The waterfall hierarchy seen through Boruvka's algorithm. A careful examination of Boruvka's algorithm for constructing MSTs give us a deeper insight in the relations between waterfalls, watersheds on graphs through invariants for the morphological opening γ_e. For the following analysis, we suppose that each node has only one lowest edge. We analyze later the general case. The waterfall hierarchy becomes apparent while constructing the MST using Boruvka's algorithm. The following different events can be distinguished by constructing iteratively a spanning tree of a graph G from an empty set of edges and nodes.

1. Two isolated nodes i and j are connected by a new edge. Since two adjacent edges have necessarily different weights, this edge is surrounded by higher edges in the graph G and hence is a regional minimum edge. Edge u is the lowest edge adjacent to i and j. Hence, as seen earlier, edge u is invariant by the opening $\gamma_e(u) = \delta_{en}\varepsilon_{ne}(u)$ and represents the lowest overflood level of the basins i and j.
2. The edge connects an isolated node i with a subtree of T. Clearly u is the lowest edge adjacent to i, otherwise node i would have been incorporated in a subtree earlier. Again edge u is invariant by the opening $\gamma_e(u) = \delta_{en}\varepsilon_{ne}(u)$.
3. The edge u connects two nodes i and j already belonging to subtrees of T. The nodes i and j have been joined to their respective subtrees through edges with lower weights than u. Hence u is not the lowest edge of the nodes i and j. For this reason u is not invariant by the opening $\gamma_e(u) = \delta_{en}\varepsilon_{ne}(u)$.

During events (1) and (2) the edges incorporated in the MST are invariant by the opening γ_e whereas events of type (3) never use edges invariant by γ_e. On the other hand, each node has one and only one lowest edge, through which it will be assigned to the MST during an event of type (1) or (2). Especially, if we only consider edges invariant by γ_e without also adding the edges of type (3), we obtain a spanning forest, and this forest is minimal as we will see in the next section. In fact this minimum spanning forest yields the same partition as the princeps waterfall algorithms proposed by Serge Beucher, in the case where G represents a region adjacency graphs, where each node represents a catchment basins and the edges are weighted with the altitude of the pass point between adjacent basins. Each basin is flooded up to the level of its lowest pass point ; like that there is no regional minimum anymore in the basins and each of them is absorbed by a neighboring basin. The edge along which this absorption takes place is precisely the lowest edge of one of its neighboring basin, that is an edge invariant by opening γ_e, belonging to the spanning forest constructed above. An example of waterfall segmentation is illustrated figure 1.

3.2 Minimum Spanning Forest

Let us now show that the spanning forest obtained by using Boruvka's algorithm is indeed minimal. If we chose an arbitrary adjacent node m_u for each minimal edge of the graph G, we obtain a family of nodes (m_u), one for each regional minimum. Adding a dummy node o linked to each m_u through an edge with valuation

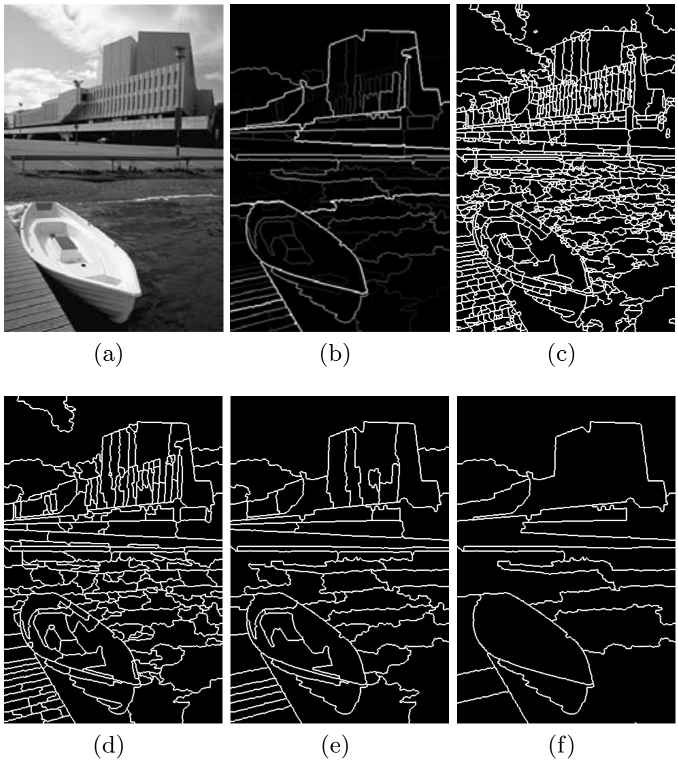

Fig. 1. (a) Original image. (b) Probabilistic gradient computed from the stochastic watershed transform [10]. (c) Watershed of the probabilistic gradient. A region adjacency graph is built from this partition, edges are weighted by the lowest passpoint between two regions. (d) Results of Boruvka's algorithm, only edges of type (1) and (2) are kept. (e-f) Second and third level of waterfall hierarchy using Boruvka's algorithm on the region adjacency graph of the partition of the previous level of the hierarchy.

-1, we get a new graph G'. Let us now construct the MST of G' starting with the dummy node o. The first steps of the algorithm visit all dummy edges with negative weights adjacent to the dummy node o. The next steps visit the edges in the same order as Boruvka's algorithm with one major difference : events of type (3) are never met, since from the beginning, there is only one tree.

After the construction is completed, we suppress the dummy node and the dummy edges and get a minimum spanning forest. This forest is identical with the forest obtained by considering only events of type (1) and (2) in Boruvka's algorithm, showing that considering only the edges invariant by opening γ_e produces indeed a minimum spanning forest. We find here by another mean the result of Serge Beucher and Beatriz Marcotegui [5], constructing the waterfall by a watershed algorithm on the MST of the neighborhood graph. This shows also that the waterfall partition of level 1 is the result of segmenting the graph G with the family of markers (m_u).

3.3 Locality or Non Locality of the Watershed ?

The edges of the spanning forest may be obtained as invariant of the opening γ_e, which is a purely local operation. Extracting the individual trees can be done via any classical labeling algorithm of connected components in a graph. It is even possible to extract a single tree, without extracting the others. This is a major difference with the watershed defined on the nodes, defined as the SKIZ of the minima for the topographic distance and constructed through competitive flooding algorithms.

4 Ambiguity in Case of Multiple Minimum Spanning Trees

In the general case, a node may have several lowest edges. Adjacent edges with the same altitude may form plateaus of any size. At each step of Boruvka's algorithm for constructing the MST, several equivalent choices of edges are possible, growing different trees. A multiplicity of spanning trees coexist, having all the same distribution of edges and producing the same nested partitions if one cuts the edges with a weight above some threshold. In particular the events of type (1), an edge connecting two nodes not connected earlier, do not necessarily produce regional minima, as this edge may belong to the inside of a plateau, which is not a regional minimum.

For this reason, one has to detect the regional minima beforehand and give the same label to all nodes belonging to the same regional minimum. After introducing as earlier a dummy node o linked by an edge of valuation -1 to each regional minimum, one may then apply Prim's algorithm for constructing the MST. Edges are considered with increasing values ; at each stage of the algorithm, only the edges adjacent to the already constructed part of the tree are considered. Nevertheless, this construction leaves a large freedom of choices and a great number of spanning trees are possible. Some of these trees do not seem desirable, as they cut the plateaus in an unfair way ; one may wish sharing the plateaus among neighboring basins by cutting them at equal distance to their lower borders.

4.1 A Hierarchical Queue Implementation for the Watershed of Edge Weighted Graphs

A fair sharing of plateaus may be reached if one resorts to the classical hierarchical queue structure: the edges are processed in increasing order, but edges with the same weight are processed according to increasing distances to the lower borders in plateaus. The structure of hierarchical queues introduces naturally the flooding order in the processing. However we loose the nice feature met above to be able to extract any tree from the watershed forest without necessarily constructing the neighboring trees which constitute its limits. In order to be able to achieve this type of extraction, one has to introduce a lexicographic order among the edges as we will see just below.

4.2 A Lower Complete Edge Graph

Our aim in this section is to complete the order relation between adjacent edges by an additional, lexicographic order relation, such that each non minimal edge has at least one neighboring edge with a lower weight. In the graph G, the edges without lower neighbor are the edges in the regional minima and the edges inside non minimal plateaus. However, each non minimal plateau has itself lower neighbors. The idea for resorbing the plateaus is to compute a geodesic distance function within the plateau towards the lower border. This distance is defined as follows : an edge u will be assigned a distance n, if n is the smallest index such $\varepsilon_E^{(n)}(u) < \varepsilon_E^{(n-1)}(u)$. This distance function may be classically obtained through a queue implementation. After this completion, each edge, except the edges inside the regional minima will have two valuations : the initial valuation $w(u)$ and the distance function $\pi(u)$ to a lower border ; the lexicographic order relation being: $u > v \Leftrightarrow \{w(u) > w(v)\}$ or $\{w(u) = w(v) \quad and \quad \pi(u) > \pi(v)\}$.

The lexicographic order relation between neighboring edges amounts to introducing a polarisation between edges. It can be represented by replacing the non oriented edges by oriented arcs. Recall that we only consider the edges invariant by the opening γ_e. Each such edge (i, j) is the lowest neighbor of one of its extremities, say i. The other extremity j is then necessarily the extremity of the lower neighboring edge of (i, j). This analysis shows that there exists an implicit orientation from lower towards higher edges ; replacing the non oriented edge (i, j) by an oriented arc $\overrightarrow{(j, i)}$ with an orientation opposite to the downwards direction conveys the same information as the lexicographic distance function of the previous paragraph. Of course the edges belonging to regional minima regions remain non oriented edges, as they do not possess a downwards direction.

With the introduction of the oriented edge graph, we will be able to extract individual trees or regions associated to a particular minimum or marker without constructing the whole watershed. It has been noticed that for the partial graph associated to the edges invariant by γ_e, neighboring basins may be connected, implying that a connected subgraph may contain various regional minima. It is not the case anymore with our oriented graph. We will associate to each minimum m_i the set of node which may be reached by an oriented path having its origin in the minimum. The ordinary Dijkstra algorithm of shortest oriented paths may be used for this extraction.

In the partial non oriented graph G', there exist nodes having two lowest edges with the same weight, leading to two distinct regional minima m_1 and m_2. Such nodes, together with all nodes for which they are downstream nodes, belong to the catchment basins of each adjacent minimum m_1 and m_2. Each of the oriented subtrees associated to these minima will contain this divide zone. Suppressing the zones which are common to two distinct neighboring oriented trees creates a minimal watershed tessellation which is not a partition, as the divide zones are missing.

4.3 Application to the Waterfall Hierarchy

The level 1 of the waterfall hierarchy consists in a minimum spanning forest, where each tree is centered on the minima. This minimum spanning forest differs from a minimum spanning tree by some missing edges. Contracting each tree of the forest to one node and introducing the missing edges produces a new tree T^1. In the case where all edges of the graph were with distinct weights, the edges of this tree also have distinct weights. And the waterfall hierarchy of level 2 is again the invariant set by opening γ_e. In all other cases, tree T^1 may have edges with the same weight and one of the general constructions above has to be applied.

5 Watershed on Node Weighted Graphs

In this last section we show how the construction of waterfalls and watersheds on edge weighted graph may be applied for constructing watersheds on node weighted graphs. In the context of segmentation, most often the watershed has to be constructed on a gradient image derived from the image. One of the problem to take into consideration is the problem of scale ; for a binary image, or a mosaic images (the image is constant on each tile of a mosaic) the gradient is a thin line and may be represented faithfully on an edge graph, in which each edge would be weighted by taking the absolute difference of the weights of its extremities. Such a local gradient does however not correctly the contours of natural images, where the boundaries of the objects are more or less blurred, corrupted by noise and the contour information is spread out on a larger surface than thin lines.

In this section we present how to transform a node weighted graph into an edge weighted graph. The watershed is then constructed on this edge weighted graph with the method presented earlier, producing a forest centered on the minima, yielding the same watershed as the watershed constructed on the node graph. The advantage of this way of doing are multiple:

- extending the waterfall hierarchy to node weighted graphs,
- transforming the problem of watershed construction into a problem of MST construction,
- after completion of the graph in order to suppress the plateaus, being able to extract individual trees, without constructing the watershed with competing markers,
- being able to construct a minimal watershed and isolate the thick divide zones.

Weighting the edges in order to stress the directions of steepest descent. Given a node weighted graph, we want to weight the edges in such a way that the watershed constructed on the edge graph is identical with the watershed on the node graphs. The weight of each edge (i, j) will be computed as follows.

1. Case where $w(i) > w(j)$. Then we compute the lower gradient $\varsigma(i) = w(i) - \varepsilon_N(i)$, obtained by subtracting from the weight of i the weight of its lowest

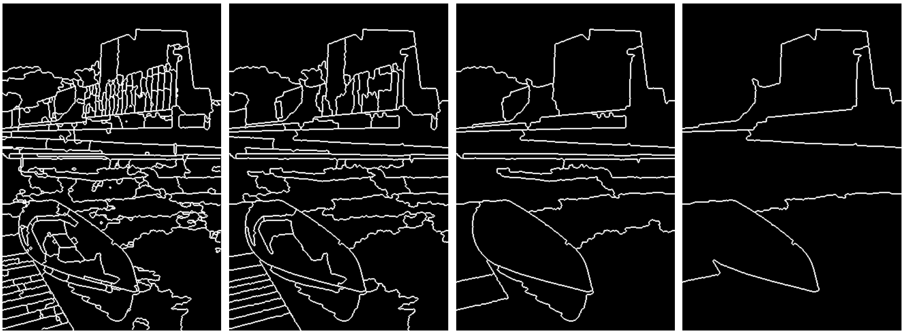

(a) (b) (c) (d)

Fig. 2. (a-d) First, second, third and forth level of waterfall hierarchy using Boruvka's algorithm on the pixel adjacency graph. The waterfall was obtained from the probabilistic gradient proposed by Angulo et al. [10] illustrated in figure 1.

neighbor. The weight of edge (i, j) is then $w(i, j) = w(j) + \varsigma(i)$. In the case where j is the lowest neighbor of i, then $\varsigma(i) = w(i) - \varepsilon_N(i) = w(i) - w(j)$ and $w(i, j) = w(i)$. In all other cases, this weight will be higher than $w(i)$.

2. Case where $w(i) = w(j)$, then we take $\min(w(j) + \varsigma(i), w(i) + \varsigma(j))$. In the case where i or j have no lower neighbor, then $\varsigma(i) = \varsigma(j) = 0$ and $w(i, j) = w(i) = w(j)$, and a plateau of edges with the same values will be created.

On this edge weighted graph, we may apply the results described earlier. A waterfall segmentation of the pixel graph obtained from a gradient image is illustrated in figure 2.

6 Conclusion

We have highlighted some new properties of the watershed and the waterfall transform through a detailed analysis of Boruvka's algorithm. These properties, linked with invariants of a specific morphological opening on graphs, provide new methods and algorithms for constructing the waterfall segmentation. We have presented these properties for both edges and nodes weighted graphs. This study brings a different point of view on the importance of minimum spanning trees for watershed and waterfall based segmentation.

References

1. Allène, C., Audibert, J., Couprie, M., Cousty, J., Keriven, R.: Some links between min-cuts, optimal spanning forests and watersheds. In: Proceedings of the The 8th International Symposium on Mathematical Morphology, Rio de Janeiro, Brazil, vol. 1, pp. 253–264 (2007)

2. Audigier, R., Lotufo, R.A.: Watershed by image foresting transform, tie-zone, and theoretical relationships with other watershed definitions. In: Proceedings of the The 8th International Symposium on Mathematical Morphology, Rio de Janeiro, Brazil, vol. 1, pp. 277–288 (2007)
3. Stawiaski, J.: Mathematical morphology and graphs: Application to interactive medical image segmentation. PhD Thesis. Mines ParisTech. (2008)
4. Meyer, F.: Grey-weighted, ultrametric and lexicographical distances. In: Computational Imaging and Vision, Mathematical Morphology: 40 Years On. Proceedings of the 7th International Symposium on Mathematical Morphology, vol. 30, pp. 289–298 (2005)
5. Marcotegui, B., Beucher, S.: Fast implementation of waterfall based on graphs. In: Ronse, C., Najman, L., Decencière, E. (eds.) Mathematical Morphology: 40 Years on: Proc. 7th ISMM, Paris, pp. 177–186. Springer, Heidelberg (2005)
6. Falcao, A., Stolfi, J., Lotufo, R.: The image foresting transform: Theory, algorithms, and applications. IEEE Transactions on Pattern Analysis and Machine Intelligence 26, 19–29 (2004)
7. Cousty, J., Bertrand, G., Najman, L., Couprie, M.: Watershed cuts: minimum spanning forests, and the drop of water principle. IEEE Trans. on Pattern Analysis and Machine Intelligence (2008) (to appear)
8. Beucher, S.: Segmentation d'images et morphologie mathématique. PhD thesis, Ecole des Mines de Paris, Paris (June 1990)
9. Beucher, S.: Watershed, hierarchical segmentation and waterfall algorithm. In: Serra, J., Soille, P. (eds.) Mathematical Morphology and its applications to signal processing (Proceedings ISMM 1994), Fontainebleau, France, pp. 69–76. Kluwer Academic Publishers, Dordrecht (1994)
10. Angulo, J., Jeulin, D.: Stochastic watershed segmentation. In: Proceedings of the The 8th International Symposium on Mathematical Morphology, Rio de Janeiro, Brazil, vol. 1, pp. 265–276 (2007)