A Qualitative Model of Gradient Flow in a Spatially Distributed Parameter

Monika Lundell

Artificial Intelligence Laboratory
Computer Science Department
Swiss Federal Institute of Technology
IN–Ecublens, 1015 Lausanne, Switzerland
Tel/Fax: +41–21–693 5209/5225
E-mail: lundell@lia.di.epfl.ch

Abstract:
This paper presents a qualitative common-sense model of gradient flow processes caused by concentration differences in a distributed parameter. A physical system is modelled by the spatial distributions of its parameters. Space is discretized into a pattern of regions according to the landmark values in the quantity space of each parameter. Gradient flow processes occur between adjacent regions with different values and create new regions of intermediate value that grow at the expense of the source and sink regions. The spatial and temporal evolution of the system depends on the relative sizes of the regions and the speeds of the flow processes. Ambiguities arise because shape information is not required. However, a plausible least complex evolution can be generated based on shape-invariant inferences and assumptions. An example of a qualitative simulation of heat flow processes is given.

1 Introduction

Models of physical systems, both quantitative and qualitative, often leave out the spatial aspects in order to keep the complexity down. These so-called lumped-parameter models describe the *temporal* evolution of parameters and are usually built on ordinary differential equations. In order to also describe the *spatial* evolution of the parameters, more complex distributed-parameter models, built on partial differential equations, are used.

In this paper, we present a qualitative distributed-parameter model of gradient diffusion, i.e. flow processes caused by concentration differences in a parameter, e.g. temperature differences. Gradient flow can be used to model many different natural phenomena, e.g. conductive heat flow, transportation of pollen and dust, water diffusion in soils, etc, as described in [Oke, 1987].

The model captures general qualitative knowledge about gradient flow processes and is suitable for situations where numerical methods cannot be used — typically in pure common-sense reasoning or when information on the spatial distribution of the parameter is sparse or unreliable. This is e.g. the case in weather forecasting when there are not enough observations to carry out a numerical simulation. The meteorologist then has to rely on common-sense reasoning and experience to produce a plausible forecast. Another example is reasoning about heat flow processes in every-day situations, like e.g. baking food in an oven or packing a grocery bag with items of different temperature.

Related work is e.g. [Collins and Forbus, 1993], where a qualitative model of thermodynamics is presented. They use a common technique in qualitative physics and model the physical system as a set of interacting objects described by parameters. The objects become the basic spatial units of the model and the variation of a parameter within an object is not described, e.g. that surrounding air will cool a hot object from the outside while the core remains hot. In our model, we take a different approach and focus on the spatial distributions of the parameters. This makes it possible to reason about the inner spatial structure of a physical object with respect to the parameter and also to disregard object boundaries when appropriate.

The model works with a qualitative view of the parameter as a pattern of discrete regions represented by the parameter's landmark values. Flow processes are triggered by features in the pattern and cause a temporal sequence of spatial modifications to the pattern. New regions and landmark values are dynamically added to the model during the simulation.

This qualitative representation of space corresponds to our assumption that information about the situation is sparse. In particular, the model does not require any knowledge about the shapes of the regions, since it is unrealistic to assume that reliable shape information could be inferred from sparse data or be present in all informal situations requiring common-sense reasoning. The exact evolution of a gradient flow process cannot be determined without shape information, as will be discussed later, but we will show how to generate a plausible least complex evolution based on shape-invariant inferences and assumptions.

In the following sections, we first describe qualitative models of spatially distributed parameters and the gradient flow processes acting on them. We then discuss the influence of gradient flow on the spatial structure of a parameter and how to handle ambiguous situations. Finally, we give an example of a qualitative simulation of heat flow processes.

2 Spatial Model

In numerical models, space, time and parameter values are continuous variables, which often leads to equations that are difficult to solve. Qualitative models attempt to overcome this problem by focusing on *landmark values*, i.e. parameter values that correspond to interesting events in the modelled system. The value domain of the parameter is modelled qualitatively as a set of landmark values, called the *quantity space*. E.g., the quantity space $\{0^\circ C\ 100^\circ C\}$ is often used to model temperature when reasoning about state changes of water (see e.g. [Forbus, 1984] and [Kuipers, 1994]). A quantity space can also consist of symbolic values, e.g. {*cold cool hot*}.

In this work, the quantity space of a parameter is used to discretize the continuous distribution of values in space into a patchwork-like pattern of adjacent regions. The initial landmark values are chosen so that all points in the space described by the parameter can be assigned such a value, either by observation, approximation or assumption. The landmark values can be either symbolic or numerical, e.g. averages of observations, as long as they provide a satisfying description of the situation as the modeller sees it. The next step is to group neighbouring points with equal landmark values into larger spatial units, i.e. *regions*. The resulting pattern is used as the qualitative spatial model.

This discrete view of space is commonly used e.g. in systems for scientific visualization (discussed in [Hagen et al., 1993]) and for certain data layers in geographic information systems (see [Laurini and Thompson, 1992]). Analyses of meteorological observations into structured weather maps (e.g. as described in [Wickham, 1970]) is another example of a discretization of space with respect to the parameters describing it.

In the spatial model, the boundary of each region in the pattern is in its turn decomposed into qualitative segments indicating its adjacency to other regions. Although the shape of a region is not assumed to be known, it is often possible to infer a relative approximation of the size of a region with respect to other regions. A quantity space for size is chosen with either symbolic or numerical landmark values.

The quantity spaces are assumed to be totally ordered and the relative magnitudes of the landmark values are supposed to be known. In case symbolic landmark values are used, a mapping to real numbers is carried out. The quantity space {*small medium extra-large*} can e.g. be mapped to {*1 3 8*} or any other sequence of magnitudes that corresponds to the modeller's view of the situation. This mapping makes it straightforward to use arithmetic operations to calculate differences and averages of landmark values and to unambiguously add new values to the quantity space.

The spatial model can be visualized as a graph where the nodes represent the regions and the edges represent the adjacencies. Figure 1 shows the temperature distribution in a field that has been unevenly heated by the sun, e.g. due to shadowing trees or partial cloudiness. The initial quantity spaces for temperature and size are shown, along with a possible mapping for size, since its quantity space contains symbolic values. The values and relative sizes of the regions are indicated by the shadowing and the sizes of the nodes respectively. In section 5, we will show a qualitative simulation of how an initial temperature distribution can be modified by heat flow processes.

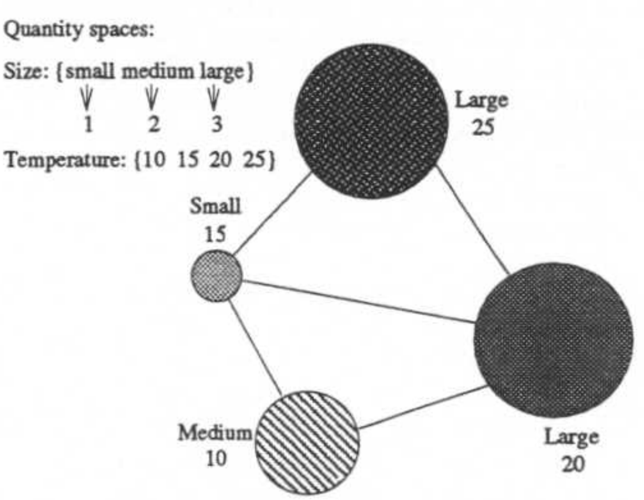

Figure 1: *Temperature distribution in an unevenly heated field.*

3 Gradient Flow Process Model

A mathematical model of gradient flow processes caused by concentration differences in a distributed parameter is the following equation (as defined in [Oke, 1987]):

$$\text{Flux Rate} = \frac{\text{Concentration Difference}}{\text{Resistance to Flow}}$$

The purpose of the flow process is to remove the concentration difference by redistributing the values of the parameter. A large concentration difference will cause a fast flow. The flow is slowed down by a large resistance parameter. The region of higher concentration is called the *source* as opposed to the *sink*. The flow will gradually move matter, e.g. heat or pollen, from the source to the sink until the concentration is homogeneous. The resulting value of the concentration depends on the sizes of the regions, i.e. how much matter they contain. A flow between two equally-sized regions will lead to an averaged concentration in both regions, i.e. one single concentration region. A flow between a small and a large region will initially only affect a part of the large region corresponding to the size of the small region. The flow will continue to move into the large region and the final result will still be one single region of homogeneous concentration, but its value will be closer to that of the large region than the small one.

The above description shows that gradient flow processes have several qualitative features. We model them with five components as follows:

Existence: A gradient flow process is triggered by the existence of two adjacent parameter regions, i.e. two connected nodes in the graph.

Effect: The gradient flow process replaces the boundary between the two regions by a growing intermediate region, called the *growth region*, that will gradually extend into both the source and the sink region, thus making them shrink. The growth will continue until at least one of the regions is completely covered, i.e. of zero size. This models the notion that volumes must match volumes, as described above. A threshold value can be defined in order to avoid infinite decomposition.

Size: The size of the growth region is initially zero and will increase at the expense of the source and the sink region.

Value: The value of the growth region is the average of the source and sink values.

Speed: A relative measure of the speed of the flow process (the flux rate) is given by the concentration difference, i.e. the value difference between the source and the sink region. We simplify the examples in this paper by assuming constant resistance. The effects of varying resistance can be modelled by intersecting the spatial distributions of the resistance and the flowing parameter, as described in [Lundell, 1994], and working with regions representing combinations of the two parameters.

Figure 2 shows the initial heat flow processes (bold dashed lines) and growth regions (tiny white circles) in the unevenly heated field in figure 1.

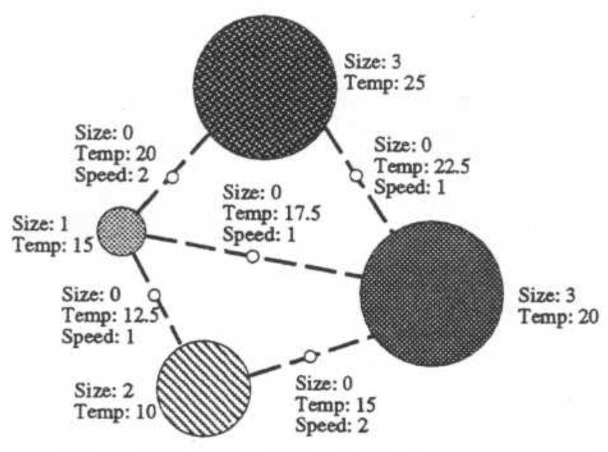

Figure 2: *Heat flow processes in an unevenly heated field.*

4 Adjacency Relations

When a gradient flow occurs, the adjacency relation between the source and the sink region is broken and replaced by a growth region adjacent to both, as shown in figure 2. As the growth region becomes larger at the expense of the source and the sink region, it will gradually take over their adjacency relations. When the source and/or sink region eventually shrinks to zero size, the growth region will have inherited all its adjacency relations. However, the exact order in which this happens can only be determined if the shapes of the regions are known. Figure 3 shows two situations that differ in the shape of the sink region but not in its adjacencies. The behaviour of the gradient flow is different as regards the adjacencies of the growth region and the connectivity of the sink region.

More complicated examples are easily constructed. We conclude that, without shape information, it is not possible to unambiguously determine the adjacency relations of the growth regions, nor the number of distinct regions. Assuming that the regions are convex or regular does not solve these ambiguities, nor is it a realistic assumption.

Shape is thus important to the evolution of gradient flow processes. However, since shape cannot be reliably inferred from sparse data and is not always known in informal situations, it is necessary to develop plausible shape-independent models of spatial processes. In this work, the evolution of the parameter regions is generated from shape-invariant inferences and assumptions are used to keep the complexity down. In this context, we define the *complexity* of a qualitative spatial situation to be the number of adjacency relations it contains, i.e. the number of potential flow processes. The *least complex evolution* is generated by updating each situation with as few new adjacency relations as possible. This technique is similar to other common simplifications in qualita-

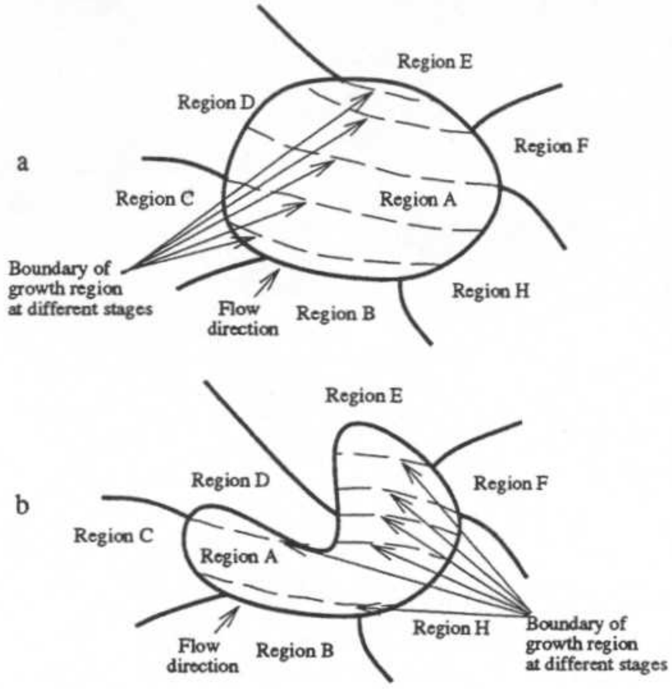

Figure 3: *In figure a, the growth region will inherit the adjacencies of region A in the following order: C/H, D, F, E. Region A remains connected during the flow. In figure b, the growth region will inherit the adjacencies in the following order: C/H, D, E, F. The flow divides region A into two disconnected parts.*

tive physics, where incompletely specified or difficult details are left out in the search for an approximate evolution that embodies the principal characteristics of the real evolution.

A common method of handling ambiguities in qualitative physics is to branch into an envisionment of all possible evolutions. In our case, this would lead to an infinite number of states, since a flow can decompose a region into an unknown number of separate pieces, as shown in figure 3. However, not all ambiguities are shape-dependent, and branching can be used when there is a finite number of choices, as will be described in the next two sections. We first discuss the shape-invariant inferences, followed by the assumptions used to generate the least complex evolution.

4.1 Shape-Invariant Inferences

The following inferences can be made regardless of shape properties:

Connectivity: The growth region is a topologically connected region, i.e. in one piece. Initially, it is also a simply-connected region, i.e. without holes, but this may change during the course of the flow. If the flow continues until the source/sink region is reduced to zero size, then that part of the growth region is again guaranteed to be a simply-connected region.

Initial adjacencies: If the source/sink region has more than one adjacency, then the growth region will inherit those that are closest to the adjacency that it is replacing. If these adjacencies also participate in flow processes, then the respective growth regions will be adjacent instead. Both cases are illustrated in figure 4.

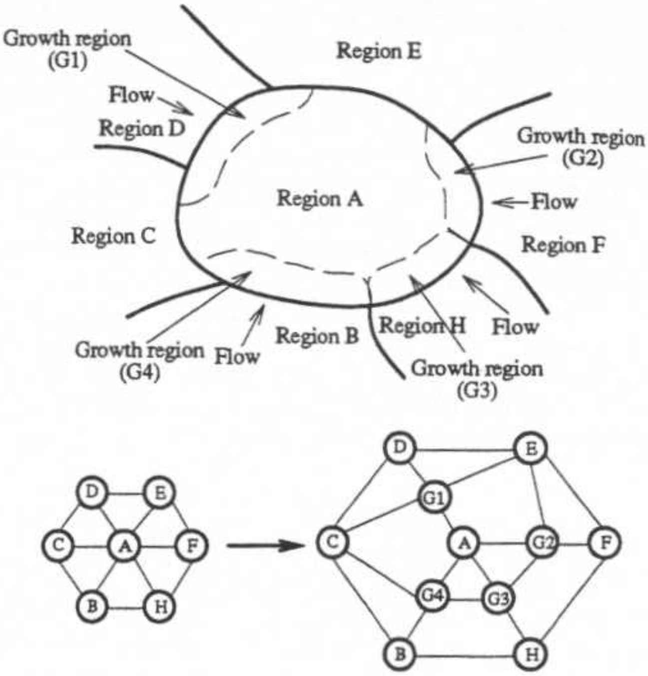

Figure 4: *Initial adjacency relations.*

Removal of regions: When the source/sink region is reduced to zero size, the growth regions that caused it to shrink will have inherited all its adjacencies. Figure 5 shows a situation where a region has six adjacencies, but only two of them participate in flow processes. It is not possible to infer with certainty which growth region will inherit which of the non-flowing adjacencies, but it is certain that the growth regions are the only candidates.

4.2 Complexity-Reducing Assumptions

The following assumptions are used to generate the least complex evolution:

Simply-connected regions: All regions are assumed to be simply-connected, i.e. without holes and in one piece, throughout the simulation. This means that no region is ever decomposed into separate parts by a flow. Without this assumption, an infinite number of states would be generated.

Growth regions meet in a point: When a source/sink region has been reduced to zero size, its

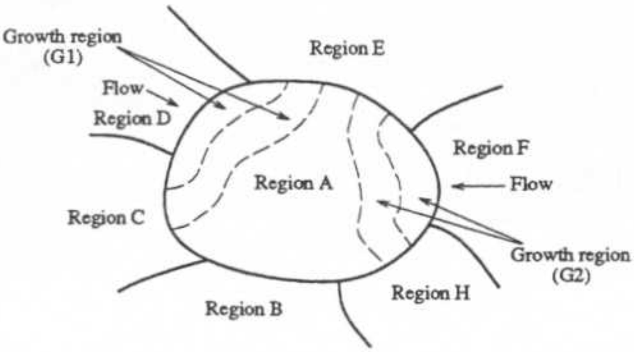

Figure 5: *Growth regions are the only candidates for remaining adjacencies.*

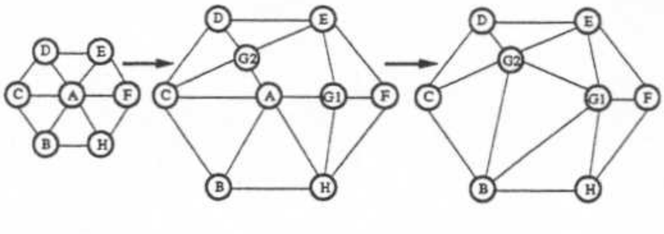

Figure 7: *Closest growth regions inherit remaining adjacencies. Graphs corresponding to the situation in figure 5.*

growth regions are assumed to meet in a point, as indicated in figure 6. This feasible evolution will not create any additional adjacencies between the growth regions, if all of the removed region's original adjacencies participated in flow processes. The alternative would be to generate all possible combinations of adjacency relations between the growth regions.

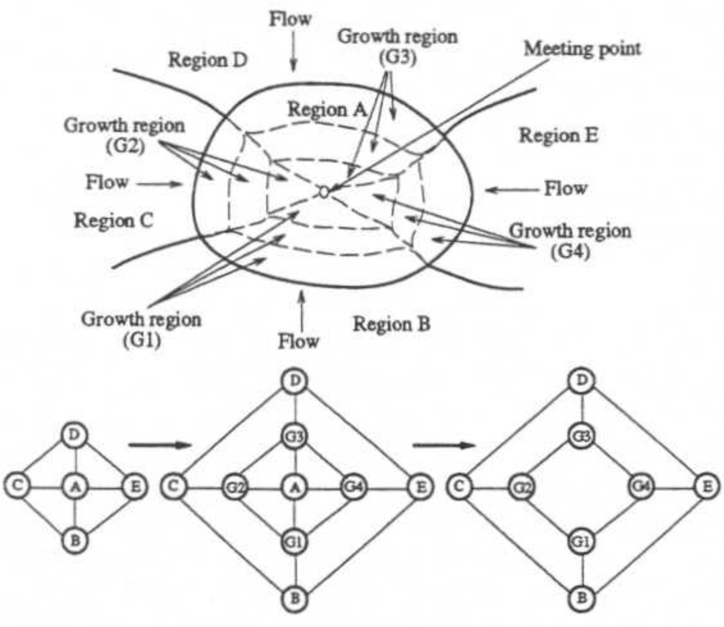

Figure 6: *Growth regions meeting in a point.*

Closest growth regions inherit remaining adjacencies: If some adjacencies remain when a source/sink region is about to be removed, as in figure 5, they are inherited by the closest growth regions. This often means that two growth regions come to share an adjacency and thus become adjacent themselves, as shown in figure 7. The alternative to this assumption would be to generate all possible combinations of growth regions and remaining adjacencies.

5 Qualitative Simulation

A qualitative simulation generates a sequence of states describing both the *temporal* and *spatial* evolution of the physical system. Each state is a graph describing a spatial situation constructed from the preceding state. The temporal evolution is described by the sequence of states. A new state is generated each time a region shrinks to zero size and is removed. This corresponds to a crucial qualitative step in the flow process model, when regions and adjacencies are created and removed. The simulation will continue as long as there are active flow processes.

The temporal order of disappearance is established by sorting the regions according to their *disappearance coefficient*. This coefficient is the ratio of the size of a region to the sum of the speeds of the influencing flow processes. The region with the smallest coefficient will be the first to disappear. Note that this is not necessarily the smallest region. The disappearance coefficient is a measure of the smallest spatial alteration, with respect to the speeds of the flow processes, that is necessary to bring about a qualitative change to the graph, i.e. terminate a flow process and cause a region to disappear.

We will explain the different steps of the algorithm with the example in figure 8. It shows a grocery bag with three adjacent items: a large packet of cold ice-cream, a medium-sized packet of refrigerated milk and a medium-sized loaf of bread at room temperature. Initially, the items correspond to distinct temperature regions. The quantity space for temperature is {*cold cool room*}, which is mapped to {*1 2 3*}, and for size {*medium large*}, mapped to {*1 2*}. These magnitudes reflect a suitable interpretation of the symbols in this particular example. In order to get a finite evolution, it is necessary to specify a threshold value that weeds out insignificant flow processes. In this example, we define a temperature difference greater than 0.5 in the mapped quantity space as significant. The regions created during the simulation will be assigned alphabetic symbols for easy reference.

The simulation algorithm proceeds in the following steps:

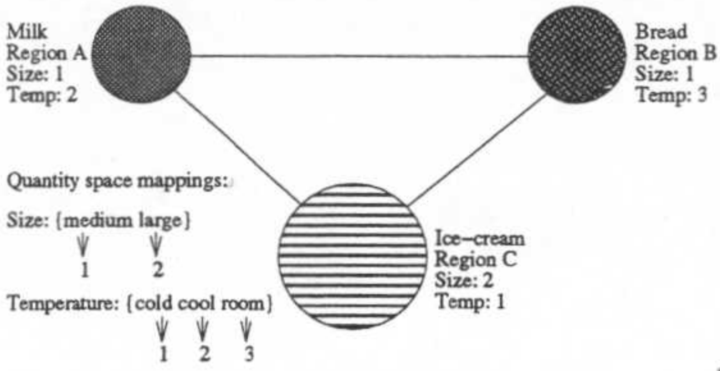

Figure 8: *Initial situation.*

Detect and install flow processes: Flow processes will occur between adjacent regions with a temperature difference greater than the chosen threshold value, i.e. between all regions in figure 8. In figure 9a, the flow processes have been installed and the adjacencies have been replaced by growth regions. Their size is initially zero and the temperature is set to the average of the source and the sink region. The speed of the flow processes is the value difference between the source and the sink region.

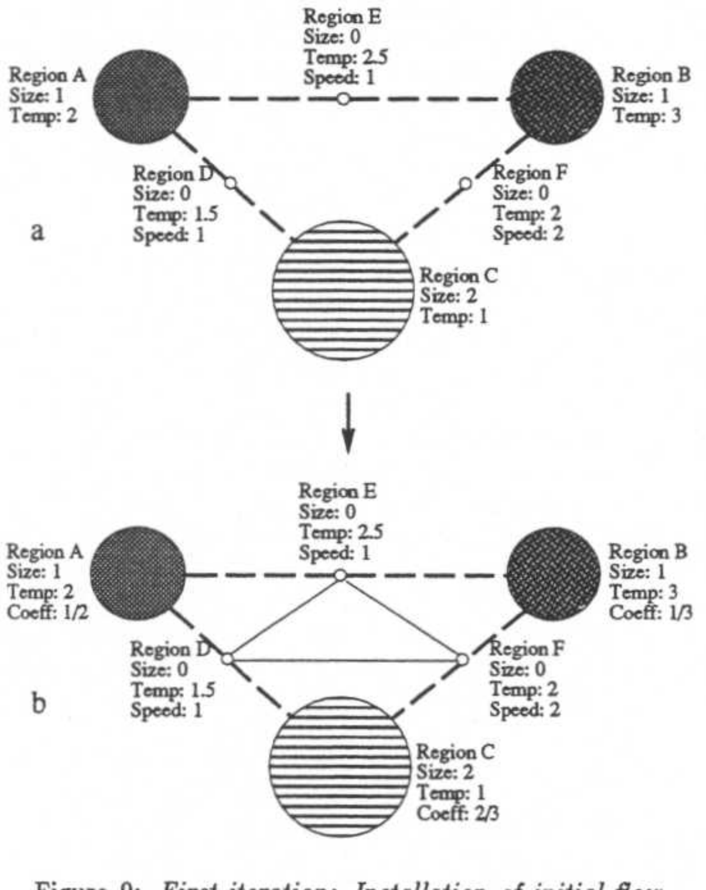

Figure 9: *First iteration: Installation of initial flow processes, adjacencies and disappearance coefficients.*

Install initial adjacencies: Once the new flow processes have been installed, initial adjacencies are installed for each new growth region, as shown in figure 9b.

Find smallest disappearance coefficient: The disappearance coefficients of the source and sink regions are calculated, as indicated in figure 9b. Region B is found to have the smallest coefficient: 1/3.

Update sizes: The speed of each flow process is multiplied by the smallest disappearance coefficient and subtracted from the sizes of the source and the sink region. It is added twice to the size of the growth region, which thus grows at the expense of the source and sink regions. The result is shown in figure 10.

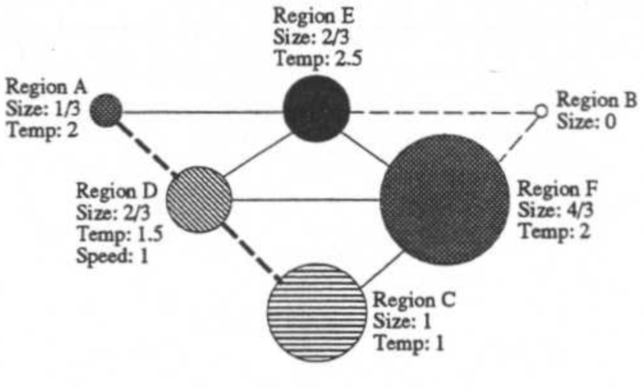

Figure 10: *First iteration: Updated sizes and removed regions.*

Terminate flow processes: The size of region B is now zero, which means the termination of the flow processes using it. The space that was originally occupied by region B, i.e. the loaf of bread, is now shared between growth regions E and F. This means that the bread now has a non-uniform temperature distribution and that its temperature as a whole has decreased. The two adjacencies of region B have been inherited by regions E and F, which are themselves adjacent. Region B and its adjacencies could be removed from the graph, but for easy comparison of subsequent graphs in this example, we only indicate its removal by replacing the adjacencies by dashed lines, as seen in figure 10.

The graph in figure 10 is the final result of the first iteration. Five regions and one active flow process remain. A second iteration is thus necessary.

The graph is reexamined and two new flow processes are detected. Figure 11a shows the new processes and the corresponding growth regions, adjacencies and disappearance coefficients. In this example, flow processes are only started between regions with a sufficiently large temperature difference, in this case regions D/E and C/F. Region D is also the growth region of the still active flow process between regions A and C. It is thus possible for a region to be involved in several processes. The simulation will determine the net influence of the processes, i.e. whether the region is growing or shrinking.

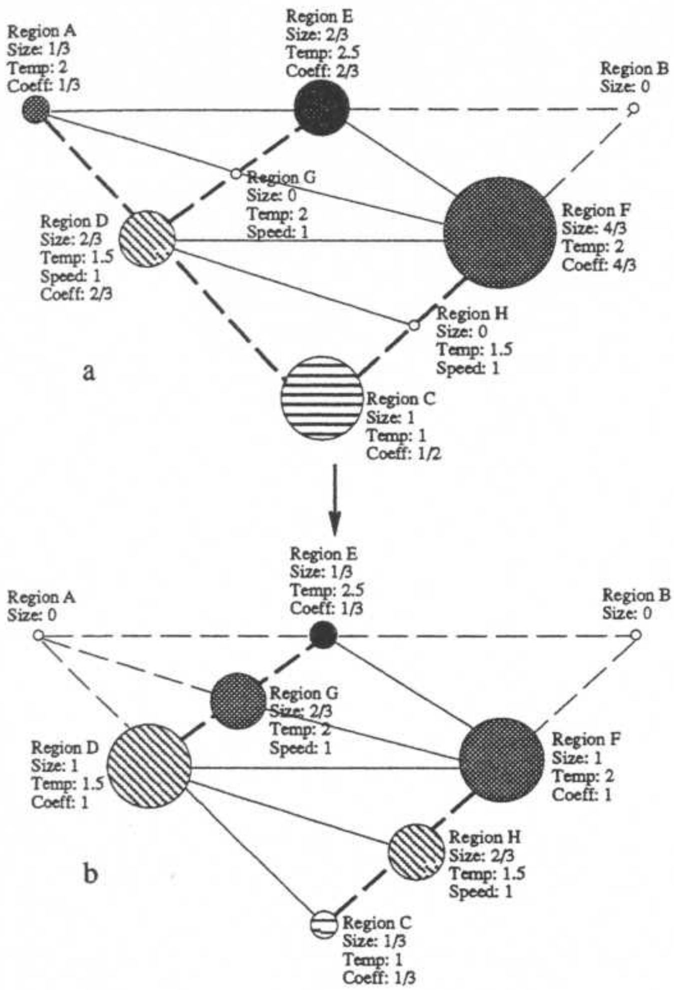

Figure 11: *Second and third iteration.*

This time, region A will be removed, since it has the smallest disappearance coefficient. The resulting graph is shown in figure 11b. Two flow processes are still active, but no new processes can be started. The minimal disappearance coefficient is found at regions C and E and both are removed from the graph, as shown in figure 12a. Four regions remain (D, F, G and H) that can be reduced to two by merging adjacent regions with equal values, as indicated in figure 12b.

The simplified graph in figure 12b is the final result of the qualitative simulation. The sum of the sizes in the final graph is the same as in the initial graph. The temperature difference between the remaining regions corresponds to the specified threshold value, indicating that the temperature is homogeneous in the system, as could be expected. The result-

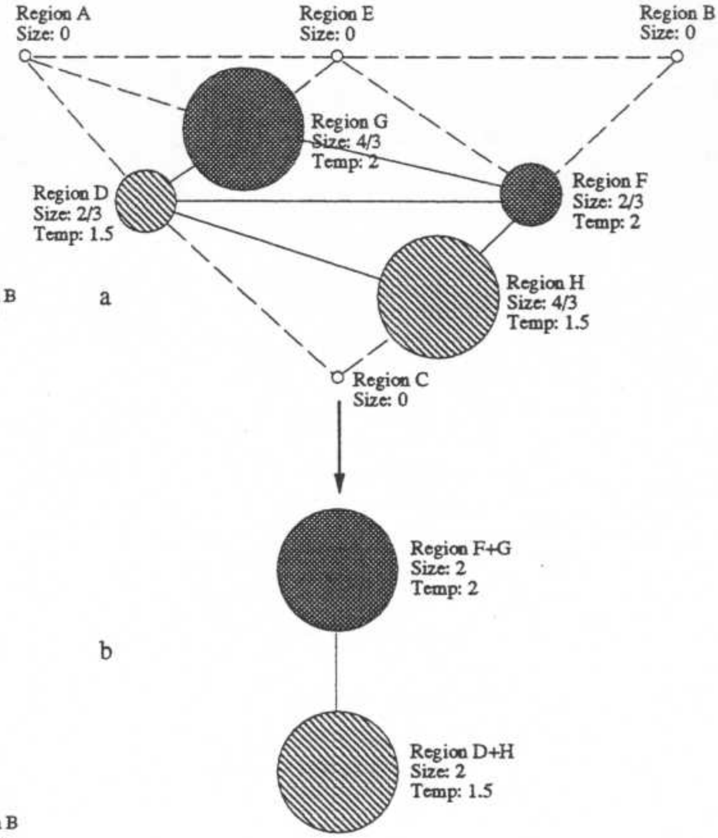

Figure 12: *Final graph and simplification: no more active flow processes.*

ing temperature lies in the interval $[1.5, 2]$, which can be mapped back to the original quantity space, giving a final qualitative temperature closer to *cool* than *cold*. The sequence of graphs produced by the simulation algorithm explains the physical background of this value.

6 Conclusion

We have presented a common-sense model and an algorithm for qualitative reasoning about gradient flow processes due to concentration differences in a distributed parameter. Since the method focuses on the spatial properties of individual parameter regions rather than objects, it is well adapted to reasoning about processes in free space and especially applications with insufficient spatial data, where the shapes of the parameter regions cannot be determined with accuracy. One example is meteorology, where the distributions of atmospheric parameters are inferred from a limited number of observation points. The shapes of the regions can in general only be roughly approximated, if at all, but it is still possible to use first principles and common-sense reasoning to draw general conclusions about the evolution of the parameter distributions.

The utility of qualitative models of spatially distributed processes is manifold. Such models would provide a user-friendly reasoning component for geographic information systems and programs for scientific visualization. They would serve as a means of communication between professionals by making it easier to share common-sense analyses of spatial situations. They would also be of great utility for pedagogical purposes.

The research presented in this paper is part of a larger project where general methods for reasoning with systems of spatially distributed parameters are being developed, see [Lundell, 1994]. The project was inspired by a study of the working methods of practising weather forecasters. We are currently developing models of basic atmospheric processes, such as radiation, conduction, convection and advection, that will be integrated in a qualitative model of a fairly complex atmospheric process: the life-cycle of a sea breeze. We are also investigating applications in agriculture and natural resource management.

7 Acknowledgements

I would like to thank the Swiss National Science Foundation for supporting this research under SPP Project No 5003-034269.

References

[Collins and Forbus, 1993] Collins, J. and Forbus, K. (1993). Building qualitative models of thermodynamic processes. Technical report, University of Illinois at Urbana-Champaign.

[Forbus, 1984] Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24.

[Hagen et al., 1993] Hagen, H., Mueller, H., and Nielson, G. M. (1993). *Focus on Scientific Visualization*. Springer-Verlag.

[Kuipers, 1994] Kuipers, B. (1994). *Qualitative Reasoning: Modelling and Simulation with Incomplete Knowledge*. MIT Press.

[Laurini and Thompson, 1992] Laurini, R. and Thompson, D. (1992). *Fundamentals of Spatial Information Systems*. Academic Press.

[Lundell, 1994] Lundell, M. (1994). Qualitative reasoning with spatially distributed parameters. In *Eighth International Workshop on Qualitative Reasoning about Physical Systems*.

[Oke, 1987] Oke, T. R. (1987). *Boundary Layer Climates*. Methuen.

[Wickham, 1970] Wickham, P. G. (1970). *The Practice of Weather Forecasting*. Meteorological Office, Her Majesty's Stationery Office.