

Storm on a Coffee Cup

Philipp K. Janert
(Dated: January 2006)

Liquid foam sometimes exhibits collective bursting of bubbles in the top layer, forming a “burst front” spreading over the foam surface. We develop a simple grid-model in an attempt to describe this behavior and discuss results and alternatives.

FOAMS

Sitting in a coffee shop, I recently noticed a remarkable phenomenon in, or rather on top of my Latte: From time to time, the bubbles in the top-most layer of foam would start bursting in one spot. The disturbance would spread, like a tiny wildfire, across part of the coffee cup, and die out eventually, exposing a fresh layer of bubbles underneath. I could observe the development of such “burst fronts” several times, starting in different locations on the foam surface, at intervals of 20 to 30 seconds. Each burst front lasted between 3 and 6 seconds.

Can we formulate a computer model to capture this behavior?

Liquid foams are strange substances. They consist to over 90% (by volume) of air or some other gas, with the remainder being formed by liquid. Yet, they can retain shape (think whipped cream) and show an elastic response — just like a solid. We encounter liquid foams in foods (e.g. whipped cream, beer or milk foam), cosmetics and detergents (soap or shaving lather), fire extinguishers, and other applications.

There are also solid foams, which are formed when the non-gas component is made out of a solid rather than a

liquid. Examples are stiff styro-foam and soft foam rubber, organic materials such as wood, bone, or bread (!), and even novel materials such as metallic foams, which combine great stiffness with very low weight.

Liquid foams are not stable, they “age”. The liquid slowly drains out of the bubble walls, and gas permeates through the walls from one bubble to neighbouring ones. Eventually, the walls become too thin and the bubbles burst, leading to the “collapse” of a foam, as every beer drinker knows.

GROWING A MODEL

When starting to formulate a model that would replicate the observed behavior, I followed two guiding principles: first of all, I wanted to make the the “ageing” of bubbles through drainage of liquid the overall driver for the dynamical evolution: at each time step, the state of *all* bubbles changes a little, leading to the popping of a bubble once its level of contained liquid falls below a critical threshold.

On the other hand, I would need to include some form of *randomness*, so as to avoid all bubbles bursting at the

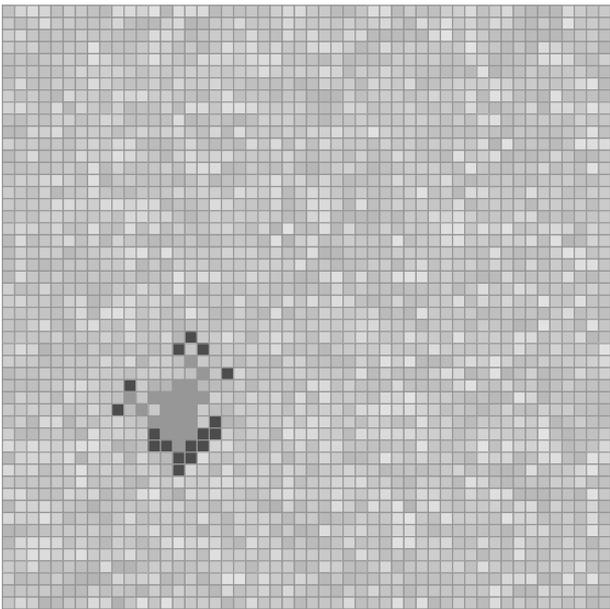


FIG. 1: A burst front starts at a seed, ...

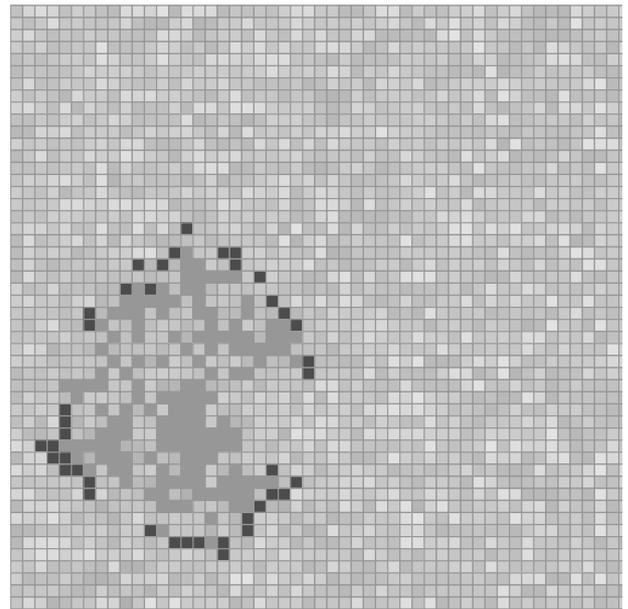


FIG. 2: ... spreads, ...

same time. Foams are inherently disordered structures, since the bubbles are not arranged in a regular fashion and have irregular, polyhedral shapes. However, after briefly looking at the algorithms for the creation of minimal lattices (using Voronoi diagrams), I decided to place the bubbles on a square grid and to capture the disorder of the foam by randomly assigning different drainage rates to each cell/bubble.

The ageing of the foam and the bursting of bubbles seem to occur on very different time scales (slow and fast, respectively), so that we can treat the bursting of bubbles as almost instantaneous, as compared to the ageing process. The program therefore operates in two different modes: all cells age collectively, but once at least one cell falls below the critical bursting threshold, the program processes only the bursting of bubbles, until the avalanche of bursting bubbles has worn itself out.

My initial implementation of these ideas did not reproduce the desired behavior. Bubbles would burst randomly all over the system, but they did not form coherent burst fronts spreading over the system.

Brief reflection revealed that I had not captured the essential physics of the problem: After all, the popping of neighboring bubbles by the initial bursting of a single bubble is a *collective* phenomenon — no bubble is an island. However, the initial model contained no spatial correlations: all bubbles were totally independent of each other.

I corrected the model in two ways: I extended the neighborhood, which would be affected through the bursting of the initial bubble. Initially, the model considered only the “nearest neighbors”, i.e. the four cells adjacent to the four edges of a square cell. I extended

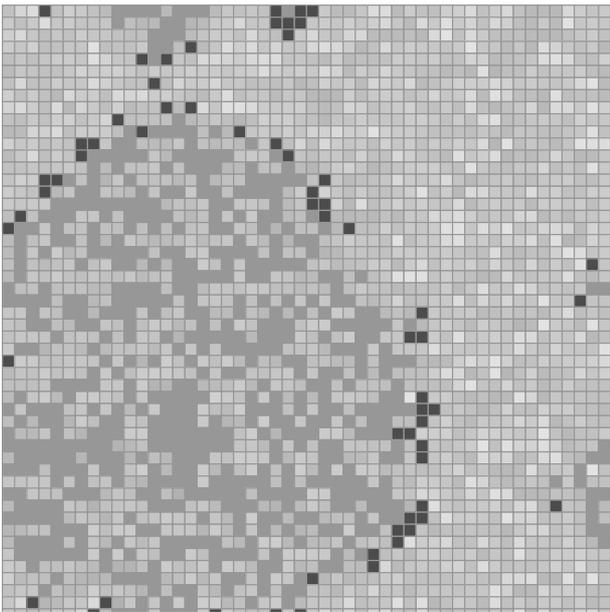


FIG. 3: ... leaving a fresh layer of foam in its wake.

this to include the “next-nearest neighbors” (those at the corners of the initial cell) and even the four primary neighbors in the next

This made the burst fronts last longer, but it did not really introduce spatial correlations into the system. To make neighboring cells resemble each other more, I introduced a step in which each cell’s (randomly assigned) drainage rate is replaced by a weighted average over its immediate neighborhood. This step can be motivated physically: if one cell drains very quickly (for whatever reason!), then the surrounding cells will also drain faster, since in addition to their own primary drainage mechanism, they can also shed liquid to the quickly draining cell in their middle.

With these modifications, the model produced reasonably realistic burst fronts. After an initial period of irregular (transient) behavior it settles down and can run for very long times.

IMPLEMENTATION NOTES

I chose Perl/Tk for the implementation of the sample program. The systems that we want to study are not very big, so that execution speed is less of a concern. However, I want to be able to follow a graphical display of the bubble system in real-time, and Tk allows me to do that in a particularly simple fashion.

Tk, as a GUI toolset, is intended to react to user events. However, it also contains ways to create synthetic events, for instance through a timer. We use the `after()` method to delay program execution at each time step. (This delay cannot be made arbitrarily small: Perl/Tk freezes if it cannot finish a task before the next occurrence of a timer event. But for slowly evolving simulations, Tk is perfectly sufficient.)

The main data structure is a two-dimensional array, with one element per lattice cell. Each element of the array is a hash, containing three keys: `fluid`, the amount of fluid contained in the cell walls, `drain`, the (random) amount of fluid that will drain per time interval, and `widget`, the ID of the `CanvasRectangle`, which represents this cell. When changing colors, we will identify each rectangle through its ID.

When the program operates in “drainage mode” (`do_drain()`), the fluid content of each cell is reduced by the cell’s (random) drainage amount at each time step. If a cell’s fluid level falls below the critical value in the `$auto_pop` variable, the cell bursts and its coordinates are added to the list `@popped` of popped cells.

During the processing of popped cells (`do_pop()`), the cells in the `@popped` list are consumed. For each cell in the list, its neighbors are examined. If their fluid level is below `$pert_pop` (for “perturbation” pop), they burst and their coordinates are added to `@popped`. The array of popped cells therefore contains the coordinates of popped cells from the previous as well as the current generation, separated by a marker (the string `'null'`).

The program also contains a routine to convert from the HSV (Hue, Saturation, Value) color model to the RGB (Red, Green, Blue) model required by Tk. I have chosen to give fresh cells a green color, which turns to yellow and finally

orange as their liquid level approaches and eventually falls below the threshold at which they can be popped by a neighbor. Cells that actually are popping are displayed in red.

Finally, the program will save screenshots (both in color and monochrome) as Postscript files upon mouse click.

QUESTIONS AND ALTERNATIVES

The greatest omission in the present model of course is the total absence of *disorder*. In a real foam, the bursting of a bubble distorts the local topology, forcing the neighboring bubbles to attempt to fill the gap. A model including not only the disordered geometry of the foam, but also the elastic behavior of the bubble walls would be quite complex, I think. The question is whether one can retain the effect of disorder in a simple lattice model.

Another question concerns the mechanism controlling the spread of the burst front. In the example program, I use one threshold of liquid content for bubbles bursting by themselves, and another, higher one for bubbles popped by their neighbors. This is a rather crude approximation. As a refinement, one could count the number of popped bubbles in the vicinity and making it ever more likely for a cell to be popped if more and more of its neighbors already burst. One can also introduce a random element into the bursting process itself.

There are many more variants one may try out. One can try a different lattice structure. A honey-comb lattice for instance looks much less artificial than the square lattice used here. One should also not forget that real foam is three-dimensional — in the present model we have assumed that all the action is happening in the top layer only, and that there is no interaction between the top- and subsequent layers. This seems like a reasonable first approximation. Can one do better?

Or one can try out different draining mechanisms. In the present model, drainage is essentially *additive*: a

fixed amount of liquid is subtracted at each time step. It might be more realistic to employ a *multiplicative* scheme, where the liquid content is multiplied by a number less than one at each time step. (In this case, all threshold values must be strictly larger than zero!) Experimentation is certainly welcome.

One question one may want to investigate with such an improved model concerns the size of each burst. Are large bursts more likely than small ones? Is there a typical size of a burst front or are big, small, and intermediate avalanches of popping bubbles equally likely? This question is left as an invitation to the reader!

FURTHER READING

There is not much literature available on the physics of foams. One recent book is **The Physics of Foams** by Denis Weaire and Stefan Hutzler (Oxford University Press, 2001). Another recent book is **Foam and Foam Films** by Dotchi Eksprova and Pyotr Kruglyakov (Elsevier, 1998).

During the preparation of this paper I discovered that the effect described here has been observed before, for instance in a series of research papers by N. Vandewalle's group at the University of Liege in Belgium: **Cascades of popping bubbles** (Physica A, 314 (2002) p. 320), **Cascades of popping bubbles along air/foam interfaces** (Physical Review E, 64 (2001), p. 021507), and references therein. These papers are mostly experimental in nature and do not attempt to formulate a model for theoretical study or simulation.

Investigation of the average size of a burst front (as suggested in the previous section) leads to the study of "Self-organized criticality". A nice introduction can be found in the book **Self-Organized Criticality** by Henrik Jeldtoft Jensen (Cambridge University Press, 1998).