A First Definition of the Blob Complex

by Aaron Mazel-Gee, lifted directly from Morrison & Walker

0 Background

Recall that given an n-dimensional system of fields and local relations $(\mathcal{F},\mathcal{U})$, we can associated to an n-manifold X its *TQFT invariant*:

$$X \rightsquigarrow A(X) = \mathcal{F}(X)/\mathcal{U}(X).$$

Here $\mathcal{U}(X)$ is the space of local relations in $\mathcal{F}(X)$, i.e. it is generated by fields on X of the form $u \bullet r$, where $u \in \mathcal{U}(B)$ is a local relation on an embedded n-ball $B \subset X$ and $r \in \mathcal{F}(X \backslash B)$. So this is just the vector space of fields on X, up to changing the field in any ball.

The blob complex can be thought of as a derived version of this construction. On a ball B, the blob complex $\mathcal{B}_*(B)$ will just be a free resolution of $H_0(\mathcal{B}_*(B))$. And when we have a short exact sequence of boundary conditions on ∂B, we obtain a long exact sequence in blob homology. In general, we will have $H_0(\mathcal{B}_*(X)) = A(X)$, and the higher blob homology groups will represent syzygies on the local relations not captured by the TQFT invariant itself.

In most of follows, we will assume that X is boundaryless. Otherwise one should fix a boundary condition in $\mathcal{F}(\partial X)$ for once and for all, and then carry out the same constructions under the assumption that everything agrees with that boundary condition.

1 An approximation of the blob complex

A blob in X is just a generalization of an embedded n-ball. In order to explain the general idea of the blob complex, for this section we will pretend that a blob really is just an embedded ball. In the next section we will explain the difference, and then the definitions given here will be very nearly correct.

Roughly, the kth level of the blob complex $\mathcal{B}_*(X)$ will be the direct sum, over configurations of k blobs $\{B_1, \ldots, B_k\}$ in X, of the vector spaces of fields splittable over those particular blobs. The boundary map $\partial : \mathcal{B}_k(X) \to \mathcal{B}_{k-1}(X)$ will just be given by (a signed sum of) erasing blobs from the picture, which certainly preserves splittability over the remaining blobs. In this section we unwind exactly what we mean for $k = 0, 1, 2$ and then return to the general case to give a full definition.

1.0 $\mathcal{B}_0(X)$

First, we define $\mathcal{B}_0(X) = \mathcal{F}(X)$.

1.1 $\mathcal{B}_1(X)$

This will basically just be $\mathcal{U}(X)$, but we want to write it out in terms of actual fields and local relations. An element of $\mathcal{U}(X)$ is just a field that splits over a blob; we'll break this down into the following data:

- a choice of a blob $B \subset X$;
- a local relation $u \in \mathcal{U}(B; c)$ (for any $c \in \mathcal{F}(\partial B)$);
- a field $r \in \mathcal{F}(X \backslash B; c)$.

Thus we define

$$\mathcal{B}_1(X) = \bigoplus_{B \subset X} \left(\bigoplus_{c \in \mathcal{F}(\partial B)} \mathcal{U}(B; c) \otimes \mathcal{F}(X \backslash B; c) \right),$$

with $\partial : \mathcal{B}_1(X) \to \mathcal{B}_0(X)$ given by $\partial(B,u,r) = u \bullet r$. So in fact $\partial(\mathcal{B}_1(X)) = \mathcal{U}(X)$, and hence

$$A(X) = \mathcal{B}_0(X)/\partial(\mathcal{B}_1(X)) = H_0(\mathcal{B}_*(X)).$$

1.2 $\mathcal{B}_2(X)$

We'll discuss $\mathcal{B}_2(X)$ before moving on to $\mathcal{B}_k(X)$, since this is the first place where we have to start worrying about how our blobs relate to each other. Here we have two blobs $B_1, B_2 \subset X$, and they will either be disjoint or nested. In both cases, we'll want to quotient $\ker(\partial : \mathcal{B}_1(X) \to \mathcal{B}_0(X))$ by fields which can be split over two blobs, identifying any two elements of $\mathcal{B}_1(X)$ that are mutually compatible.

For two disjoint blobs $B_1, B_2 \subset X$ (actually they only need disjoint interiors), we're looking for fields $x \in \mathcal{F}(X)$ from which we can simultaneously split off local relations on B_1 and B_2, i.e. $x = u_1 \bullet u_2 \bullet r$ for some $r \in \mathcal{F}(X \setminus (B_1 \cup B_2))$. So we make a definition similar to the one above, but we don't actually want to keep track of the ordering of our blobs so we set

$$\mathcal{B}_{2,disjoint}(X) = \left(\bigoplus_{\substack{B_1,B_2 \subset X \\ B_1^\circ \cap B_2^\circ = \emptyset}} \bigoplus_{\substack{c_1 \in \mathcal{F}(\partial B_1) \\ c_2 \in \mathcal{F}(\partial B_2)}} \mathcal{U}(B_1; c_1) \otimes \mathcal{U}(B_2; c_2) \otimes \mathcal{F}(X \setminus (B_1 \cup B_2); c_1, c_2) \right) \Bigg/ \left\{ \begin{array}{l} (B_1, B_2, u_1, u_2, r) = \\ -(B_2, B_1, u_2, u_1, r) \end{array} \right\}.$$

For two nested (possibly equal) blobs $B_1 \subset B_2 \subset X$, we're looking for fields $x \in \mathcal{F}(X)$ from which we can split off a local relation on B_1 or B_2. But note that we actually only need to assume we can split off a field over B_2, since local relations are contagious. So we define

$$\mathcal{B}_{2,nested}(X) = \bigoplus_{B_1 \subset B_2 \subset X} \bigoplus_{\substack{c_1 \in \mathcal{F}(\partial B_1) \\ c_2 \in \mathcal{F}(\partial B_2)}} \mathcal{U}(B_1; c_1) \otimes \mathcal{F}(B_2 \setminus B_1; c_1, c_2) \otimes \mathcal{F}(X \setminus B_2; c_2).$$

Now that we have exhausted all types of blob configurations we define $\mathcal{B}_2(X) = \mathcal{B}_{2,disjoint}(X) \oplus \mathcal{B}_{2,nested}(X)$, and (for either type of configuration) we define $\partial(B_1, B_2, u, r', r) = (B_2, u \bullet r', r) - (B_1, u, r' \bullet r)$.

1.3 $\mathcal{B}_k(X)$

Hopefully the pattern should be clear by now. We define $\mathcal{B}_k(X)$ to be generated by fields on x along with configurations of k blobs that are all disjoint or nested and over which x splits, such that if a blob B_i does not strictly contain any other blob (this is called a *twig*) then x must restrict to a local relation in $\mathcal{U}(B_i)$. (Again, this last condition ensures that x restricts to a local relation on any blob, by the condition that local relations form an ideal.) In symbols, if we let

- $T \subseteq \{1, \ldots, k\}$ be the indices of the twig blobs (for some particular blob configuration)
- $X' = X \setminus (B_1 \cup \ldots \cup B_k)$
- $c \in \mathcal{F}(\partial X')$ be the sum of the boundary conditions $c_i \in \mathcal{F}(\partial B_i)$

then

$$\mathcal{B}_k(X) = \left(\bigoplus_{\substack{B_1,\ldots,B_k \subset X \\ B_i, B_j \text{ disjoint} \\ \text{or nested}}} \bigoplus_{c \in \mathcal{F}(\partial X')} \left(\bigotimes_{i \in T} \mathcal{U}(B_i; c_i) \right) \bigotimes \left(\bigotimes_{i \notin T} \mathcal{F}(B_i \setminus \bigcup_{j \neq i} B_j; c_i, c_j) \right) \bigotimes \mathcal{F}(X'; c) \right) \Bigg/ \begin{array}{l} \text{permutations} \\ \text{of blobs} \\ \text{with signs.} \end{array}$$

The boundary map $\partial : \mathcal{B}_k(X) \to \mathcal{B}_{k-1}(X)$, as stated before, is the alternating sum of erasing one of the balls.

2 Blobs

The following definitions are motivated by the fact that we would like the following two operations on blob configurations to yield blob configurations:

- For any (possibly empty) blob configuration $\{B_1, \ldots, B_k\}$ on an n-ball B, $\{B_1, \ldots, B_k, B\}$ is also a blob configuration (i.e. we can always add B itself as an outermost blob).

- If we obtain X_{gl} from X by gluing, then any blob configuration on X gives a blob configuration on X_{gl}.

However, allowing these operations gives blob configurations whose complements are not manifolds. For example, suppose we have two $1 \times 1 \times 2$ blocks that have each been decomposed into two blobs as in the picture, and suppose we are planning to glue these two blocks together in the evident way.

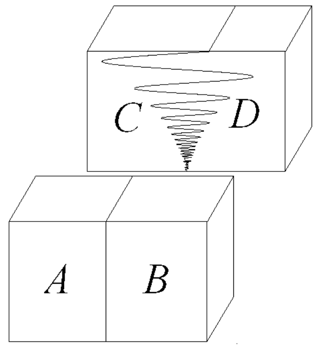

Certainly we must allow $\{A\}$ as a blob configuration in $A \cup B$ and $\{D\}$ as a blob configuration in $C \cup D$. But then we also must allow $\{A, D\}$ as a blob configuration in $(A \cup B) \cup_{\text{face}} (C \cup D)$, whose complement is not a manifold. This motivates the following sequence of definitions.

Definition. We define a *gluing decomposition* of a manifold X to be a sequence of manifolds $M_0 \to M_1 \to \ldots \to M_m = X$ such that each M_k is obtained from M_{k-1} by gluing together a disjoint pair of homeomorphic $(n-1)$-manifolds in the boundary of M_{k-1}. If M_0 is a disjoint union of balls, then the gluing decomposition is a *ball decomposition*. We say that a field on X is *splittable* along the gluing decomposition $M_0 \to \ldots \to M_m = X$, it is the image of a field on M_0.

So all the points of X are already contained in M_0, and we just need to find a sequence of gluings such that at every stage we still have a manifold.

Definition. A *blob configuration* in X is an ordered collection of k subsets $\{B_1, \ldots, B_k\}$ of X such that there exists a gluing decomposition $M_0 \to \ldots \to M_m = X$ for which each B_i is the image in X of some connected component M'_j of some M_j, where M'_j must be a ball. We say that the gluing decomposition is *compatible* with the configuration.

Thus, the final $1 \times 2 \times 2$ block in the above example can be realized as the ball decomposition

$$A \sqcup B \sqcup C \sqcup D \to (A \cup B) \sqcup (C \cup D) \to A \cup B \cup C \cup D,$$

and $\{A, D\}$ is indeed a legal blob configuration. Observe the following general facts:

- Any two blobs must be nested or have disjoint interiors.
- Nested blobs may have boundaries that overlap (or even coincide).
- Blobs may meet ∂X.
- Through the sequence of gluings, M'_j may have been glued to itself, and so blobs need not actually be embedded balls.
- Complements of blob configurations need not be manifolds.

Recall that our old definition of the blob complex involved choosing fields on the complements of blob configurations. But we can only choose fields on manifolds. In light of the final observation above, we make a new definition.

Definition. A *k-blob diagram* on X is a configuration of k blobs $\{B_1, \ldots, B_k\}$ on X and a field $r \in \mathcal{F}(X)$ which is splittable along some gluing decomposition compatible with the configuration, such that if the twig B_i is the image of the ball M'_j then the restriction of r to M'_j is a local relation.

(Remember that we can only define local relations on balls.)

When we write $\{B_1, \ldots, B_k\}$ we mean a configuration of k blobs, and when we write $(\{B_1, \ldots, B_k\}, r)$ we mean a k-blob diagram. Now we can succinctly define $\mathcal{B}_*(X)$ by

$$\begin{aligned} \mathcal{B}_k(X) &= \left(\bigoplus_{\{B_1,\ldots,B_k\}} \{(\{B_1, \ldots, B_k\}, r)\} \right) \Big/ \begin{array}{l}\text{permutations}\\ \text{of blobs}\\ \text{with signs}\end{array} \\ \partial(\{B_1, \ldots, B_k\}, r) &= \sum_{i=1}^{k} (-1)^{i+1} (\{B_1, \ldots, \hat{B}_i, \ldots, B_k\}, r). \end{aligned}$$

It is immediate that ∂ really does take $\mathcal{B}_k(X)$ into $\mathcal{B}_{k-1}(X)$.

3 The blob laws

Throughout Morrison & Walker's original paper, results are mostly proved using the blob laws rather than the actual definition of blob homology. These blob laws conjecturally characterize the blob complex, but this has not yet been established.

1. *The blob complex is functorial with respect to homeomorphisms.*

2. *There is a natural isomorphism* $\mathcal{B}_*(X \sqcup Y) \cong B_*(X) \otimes B_*(Y)$ *(with the usual convention* $\partial(a \otimes b) = (\partial a) \otimes b + (-1)^{|a|} a \otimes (\partial b)$*).*

   Proof. We can combine any pair of blob diagrams on X and Y to a blob diagram on $X \sqcup Y$ by listing first the blobs on X, then the blobs on Y. Up to sign, every blob diagram on $X \sqcup Y$ arises in this way, and this association is compatible with the boundary maps because of the signs chosen in their definitions. □

3. *Let* $c \in \mathcal{F}(\partial B)$ *be any boundary condition. If the natural quotient map* $p : \mathcal{B}_*(B; c) \to H_0(\mathcal{B}_*(B; c))$ *has a splitting* $s : H_0(\mathcal{B}_*(B; c)) \to \mathcal{B}_0(B; c)$*, then these two maps induce a chain homotopy equivalence between* $\mathcal{B}_*(B; c)$ *and the complex* $H_0 = \cdots \to 0 \to H_0(\mathcal{B}_*(B; c)) \to 0 \to \cdots$.

   Proof. By assumption $ps = \mathrm{id}_{H_0}$, so we just need a collection of maps $h : \mathcal{B}_k(B; c) \to \mathcal{B}_{k+1}(B; c)$ such that $\partial h + h\partial = \mathrm{id}_{\mathcal{B}_*(B;c)} - sp$. For $k \geq 1$ we define $h_k((\{B_1, \ldots, B_k\}, r)) = (\{B_1, \ldots, B_k, B\}, r)$, and we define $h_0(r) = (\{B\}, r - s(p(r)))$. This gives the diagram

$$\begin{array}{ccccccccc} \cdots \longrightarrow & \mathcal{B}_3(B;c) & \longrightarrow & \mathcal{B}_2(B;c) & \longrightarrow & \mathcal{B}_1(B;c) & \longrightarrow & \mathcal{B}_0(B;c) & \longrightarrow 0 \\ & \downarrow{\scriptstyle \mathrm{id}} & \swarrow{\scriptstyle h_2} & \downarrow{\scriptstyle \mathrm{id}} & \swarrow{\scriptstyle h_1} & \downarrow{\scriptstyle \mathrm{id}} & \swarrow{\scriptstyle h_0} & \downarrow{\scriptstyle \mathrm{id}-sp} & \\ \cdots \longrightarrow & \mathcal{B}_3(B;c) & \longrightarrow & \mathcal{B}_2(B;c) & \longrightarrow & \mathcal{B}_1(B;c) & \longrightarrow & \mathcal{B}_0(B;c) & \longrightarrow 0. \end{array}$$

   It is obvious that $\partial h + h\partial = \mathrm{id} - sp$ on $\mathcal{B}_k(B; c)$ for all $k \neq 1$. At $k = 1$ we have

$$\begin{aligned} (\partial h_1 + h_0 \partial)(\{B_1\}, r) &= \partial((\{B_1, B\}, r)) + h_0(r) \\ &= (\{B_1\}, r) - (\{B\}, r) + (\{B\}, r - s(p(r))) \\ &= (\{B_1\}, r) + (\{B\}, -s(p(r))). \end{aligned}$$

   But note that $p(r) = 0$ by definition of blob homology, so in fact this map is the identity on $\mathcal{B}_1(B; c)$. □

   Assuming $(\mathcal{F}, \mathcal{U})$ is enriched over **Vect**, we will always have such a splitting. But note that even when there is no such splitting, we can still let $h_0 = 0$ and get a homotopy equivalence between $\mathcal{B}_*(B; c)$ and $\cdots \to 0 \to \mathcal{U}(B; c) \to \mathcal{F}(B; c)) \to 0 \to \cdots$.

4. *If* X *is a disjoint union of balls, then* $\mathcal{B}_*(X; c)$ *is contractible.*

   Proof. This follows directly from Properties 2 and 3. □

5. *Suppose $\partial X = Y \cup Y \cup Z$. Let X_{gl} be the result of gluing the two copies of Y together, and write $\partial X_{\mathrm{gl}} = Z_{\mathrm{gl}}$. Suppose $c \in \mathcal{F}(X)$ restricts to the same boundary condition $a \in \mathcal{F}(Y)$ on both copies of Y. Then we can necessarily glue the restriction $b \in \mathcal{F}(Z)$ of c to itself to get $b_{\mathrm{gl}} \in \mathcal{F}(Z_{\mathrm{gl}})$. For any such situation, there is a chain map*

$$\mathrm{gl} : \mathcal{B}_*(X; a, a, b) \to \mathcal{B}_*(X_{\mathrm{gl}}; b_{\mathrm{gl}})$$

which is natural with respect to the actions of diffeomorphisms and iterated gluings.

4 A combinatorial aside

Blob configurations are rather combinatorial in nature. In this section we describe a functorial, simplicial set-like construction which associates to any blob configuration what we will call a *cone-product polyhedron*. We denote by $\mathcal{P}$ the collection of these, and we denote a typical object by the letter ρ.

Note that this construction is ignorant of whether nested blobs have intersecting boundary.

From a blob configuration b we build a simplicial complex $p(b)$ as follows:

- Let $p(\emptyset) = \mathrm{pt}$, where $\emptyset$ denotes a 0-blob diagram.
- If b and b' are non-overlapping blob diagrams (i.e. the interiors of their blobs are disjoint), let $p(b \sqcup b') = p(b) \times p(b')$ (note that this rule makes the previous rule acceptable).
- If $\bar{b}$ is obtained from b by adding an outer blob which encloses all the others, let $p(\bar{b}) = \mathrm{cone}(p(b))$.

Thus, (assuming our diagram has any blobs at all) we start with an edge for each twig blob, take a Cartesian product whenever we need to combine two configurations, and take the cone whenever a new blob encloses the existing configuration. So for example, a diagram of k nested blobs yields a k-simplex, while a diagram of k disjoint blobs yields a k-cube. If two or more blobs are equal, this is still well-defined if we consider them as being nested in any arbitrary way.

We can then strengthen the construction to commute with taking the boundary, but we must make a new definition of “boundary” for our cone-product polyhedra. Let $\Phi\mathcal{P}$ denote the graded free abelian group on $\mathcal{P}$, with $\Phi\mathcal{P}_k$ generated by cone-product polyhedra arising from k-blob configurations. Note that throughout the construction of a cone-product polyhedron we can keep track of a distinguished point $*$, namely the point associated to the 0-blob diagram. We define a homomorphism $\delta : \Phi\mathcal{P}_k \to \phi\mathcal{P}_{k-1}$ by, for any $\rho \in \mathcal{P}$, setting $\delta(\rho)$ to be a signed sum of those faces in $\partial\rho$ which contain $*$, except that we define $\delta(*) = *$. (Thus $*$ may be called the *permanent vertex*.) This gives us (up to signs that have yet to be checked) that $p(\partial b) = \delta p(b)$, since this is true when we modify b and $p(b)$ using our three operations:

- If $b = \emptyset$ is the 0-blob diagram, then

$$\begin{aligned} p(\partial\emptyset) &= p(\emptyset) = * \\ \delta(p(\emptyset)) &= \delta(*) = *. \end{aligned}$$

- If b and b' are non-overlapping blob diagrams, then

$$\begin{aligned} p(\partial(b \sqcup b')) &= p((\partial b) \sqcup b' + (-1)^{|b|} b \sqcup (\partial b')) = p(\partial b) \times b' + (-1)^{|b|} p(b) \times p(\partial b') \\ \delta p(b \sqcup b') &= \delta(p(b) \times p(b')) = \delta(p(b)) \times p(b') \pm p(b) \times \delta(p(b')) \end{aligned}$$

  so the statement follows by induction on $|b|$ and $|b'|$.

- If $|b| = k$, then

$$\begin{aligned} p(\partial(\bar{b})) &= p\left(\sum_{i=1}^{k} (-1)^{i+1} \overline{b_{\hat{i}}} + (-1)^k b\right) = \sum_{i=1}^{k} (-1)^{i+1} \mathrm{cone}(p(b_{\hat{i}})) + (-1)^k p(b) \\ \delta p(\bar{b}) &= \delta(\mathrm{cone}(p(b))) = \sum_{i=1}^{k} \pm \mathrm{cone}((p(b_{\hat{i}}))) \pm p(b). \end{aligned}$$

In other words, $p : (\mathcal{B}_*, \partial) \to (\Sigma\mathcal{P}, \delta)$ is a homomorphism of differential graded groups.

Presumably there should be a way to carry over the information on blobs to $\Phi\mathcal{P}$ in such a way that we can compute $H_*(\mathcal{B}_*(X))$ from its image in (a souped-up version of) $\Phi\mathcal{P}$. This might involve something like a sheaf over the configuration space $\mathrm{Conf}_{\mathcal{B}}(X)$ of blobs on X of fields on X that determine allowable k-blob diagrams. This would bring to bear the existing machinery of simplicial objects, as we might hope to keep track of the fields, local relations, and coherence data on our cone-product polyhedra. However, the ambiguity in their construction when we have equal blobs might present problems, although it should be noted that if $B_i = B_j$ then $\mathcal{F}(B_i \backslash B_j) = \mathcal{F}(B_j \backslash B_i) = k$ (the base field) by blob law 2, so perhaps this does not immediately crush all hope for such an approach.