

Chapter 1

Parallel Boost Graph Library

Contributed by Trenton W. Ford

1.1 Introduction

The Parallel Boost Graph Library(**PBGL**) is a part of the Boost Graph Library(**BGL**). The parallel libraries offers packages that focus on the distribution of storage and computation of graphs and graph algorithms. The PBGL uses the MPI scheme Since 1999 the BGL has been in the Boost C++ Libraries commonly referred to as Boost[1, 2]. Since 2008 the PBGL has been part of BGL. Boost is a set of more than 80 separate libraries designed to be fast, relatively lightweight, and as generic as possible. Similarly, these same library implementation paradigms and techniques define the PBGL.

For instructions on how to install and link the Boost and the Parallel Boost Graph Libraries, please visit the Boost Library Documentation for the current release version¹.

1.2 Background

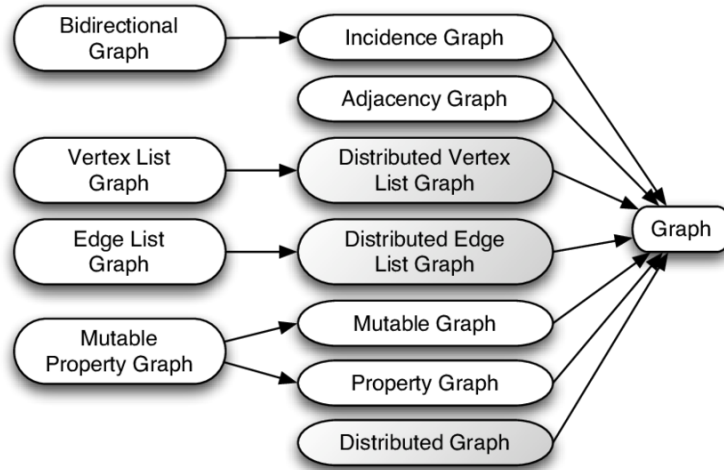
Discuss what was the reasons for the paradigm, what was it derived from, what kind of tool chain it is, what in general is its objective, where can you get code.

Within the PBGL, the central theme of design was genericism. The decision to make a maximally generic graph library was driven the developers' realization that there are many graphs with different properties on top of which many algorithms can operate. These algorithms have associated properties and underlying structures as well. With this understanding it became clear that for any C++ based graph library to have broad applicability, it would require a well-defined graph structure segmentation (underlying data-structure, functional interface, graph object properties, etc.). This structure will be described in detail in Section 1.4.

Prior to BGL beginning development, there was already strong movement towards the creation of maximally generalizable code (*generic programming*) libraries in C++ due to the addition of templating and the Standard Template Library(**STL**) in 1994 and the need for better code reuse within the language. The BGL was developed in response to the growing applicability of graphs as convenient computational representations. The need for a subset of the BGL to be dedicated to the parallelization of the ideas and concepts distilled during the BGL's development became more evident as larger and more complex graphs were beginning to gain popularity and utility.

¹<https://www.boost.org/doc/>

Figure 1.1: PBGL specific concepts are shown in gray, while BGL concepts otherwise.



1.3 Expressing Graphs

Within the BGL, graphs can be stored in two main structural types: adjacency lists and adjacency matrices. The PBGL extends these storage methods to handle distributed storage of the data-structures while maintaining the same access and manipulation methods as the BGL exposed. This makes porting code written to work with the BGL straightforward to parallelize using the PBGL. The BGL also includes methods to use other storage schema while still exposing the same standardized and generic BGL interface for developers. This makes the BGL, and by extension, the PBGL quite adept at working in congruence with relational database platforms among other software. To fully understand how graphs are defined within the library, it is first worth understanding how a graph is deconstructed as it relates to the design of the graph generics. These ideas are encapsulated in Section 1.3.1.

1.3.1 Graph Concepts

A graph concept, as it relates to the BGL, is the set of syntactic requirements, function prototypes, as well as semantic requirements that are necessary to make the BGL as generic as possible, while still maintaining the ability to be flexible upon full instantiation of the template. The idea of concepts is a crossover from general generic programming paradigms within C++. For a more through introduction to these concepts a great resource is the book *Generic Programming and the STL*.

A large component of the PBGL design process depends on the BGL upon which it was built. From the BGL, the PBGL adopts its usage of the graph Concepts Taxonomies (*fig 1.1*), but extends it to capture distributed graph concepts such as graphs built on distributed adjacency lists or distributed adjacency matrices.

The concepts themselves could be more easily considered as taxonomy of graph types where the separations of graph types is driven by the barriers imposed by generic programming and the strongly typed nature of C++. Said another way, the taxonomies served as natural delineations between the requirements and constraints of each graph type.

For a more in-depth look at the rationale behind the graph concepts, please see Douglas and Lumsdaine's original PBGL paper[2].

1.4 Syntax

The Boost libraries have a reputation for being syntactically difficult. This perception is driven mainly by the nature of templating and generic programming in C++. For a developer interested in using Boost and the PBGL it is often useful to clearly define the graph types that the application calls for, and to either find or write a wrapper class that implements the appropriate graph type and exposes an interface that is more convenient. For simple enough graphs using the libraries unwrapped is relatively simple. Sections 1.4.1-1.4.3 contain simple use cases and some example code.

1.4.1 Simple Graph Creation

The most useful introduction to the PBGL is through a simple example or creating a graph. When deciding on the set of graph concepts to utilize, it's usually prudent to first consider the use-case for the graph. For instance, the following graph could pattern a family tree structure where we imagine that the edges are directed and the sparcity of the graph structure makes an adjacency list the optimal underlying data-structure.

Listing 1.1: Create a Simple Family Tree Graph

```

1 #include <boost/graph/adjacency_list.hpp> // for customizable graphs
2 // Construct a graph with the vertices container as a vector
3 enum family {Bob, Tod, Rob, Jeb, Sue, Ann, Mae, Rea, N}
4 void main()
5 {
6     typedef boost::adjacency_list<boost::vecS, boost::vecS, boost::directedS> Graph;
7     // graph class is adjacency_list, vertices and edges will be stored in vectors
8     // The graph will be undirected
9     Graph g(8); //Create a graph sized for 8 vertices
10
11     add_edge(Bob, Sue, g);
12     add_edge(Bob, Mae, g);
13     add_edge(Bob, Rob, g);
14     add_edge(Sue, Ann, g);
15     add_edge(Mae, Tod, g);
16     add_edge(Rob, Jeb, g);
17     add_edge(Rob, Rae, g);
18 }
```

The above family tree example contained no parallel components, and relied fully on the serial BGL, but this example serves the purpose of introducing the syntax of the generic library.

1.4.2 Property Maps

The PBGL uses property maps in multiple incarnations to serve as convenient associative methods for graph attributes. Most prominently property maps can be used to relate properties or attributes to *vertex descriptors* and *edge descriptors*. Here is an example of a vertex property and an edge property:

```

1 struct VertexProperties
2 {
```

```

3  std::string name;
4  }
5
6  struct EdgeProperties
7  {
8      double weight;
9  }

```

If we wanted to map these properties to graph objects, the BGL offers three functions that facilitate the property map associations: *get(p-map, key)*, *put(p-map, key, value)*, and *p-map[key]*. Here's an example of using the get, set, and key function:

```

1  \\ Create a vector map using the provided boost type
2  boost::vector_property_map<VertexProperties , indexMapType>
3  dataMap(num_vertices(graph), indexMap);
4
5  \\ Select a random vertex from the graph using an RNG
6  Vertex v = random_vertex(graph, rng);
7
8  \\ Get the vertex properties from vertex v
9  auto vp = get(dataMap, v);
10
11 \\ Change the values in the property
12 vp.set_current_values(1,1,0);
13
14 \\ Associate the newly changed vertex properties with vertex v
15 put(dataMap, init_infected, vp);

```

1.4.3 Graph Traversal

Moving from one vertex in a graph to another across an edge is called graph traversal. For example, if in the family tree example you wanted to find one of Bob's relatives, you could start from the Bob vertex and traverse one move/step/hop to one of Bob's relatives. Vertices that are connected by an edge are called adjacent vertices. This idea of graph traversal is a foundational activity in graphs and especially graph algorithms. This activity exists in both directed and undirected graphs.

Traversing a graph in the BGL is relatively simple as the activity is facilitated through the use of generic iterator classes. There are five types of iterator generics used for traversal: *vertex iterator*, *edge iterator*, *out-edge iterator*, *in-edge iterator*, and *adjacency iterators*.

```

1  Vertex Step(Vertex current_node, Graph G) {
2
3      double random_number = ud(gen);
4
5      // Create two out_edge_iterators of the same type "Graph" as graph G
6      graph_traits<Graph::GraphContainer>::out_edge_iterator eit, eend;
7
8      // Select all outgoing edges from current_node
9      std::tie(eit, eend) = out_edges( current_node, G);
10
11     for( auto it = eit; it != eend; ++it) {
12         //loop through all the edges and store the last seen target vertex

```

```


13     Vertex target = boost::target(*it, G);
14     }
15     // return the last seen target vertex
16     return target;
17
18 }
```

1.5 Key Graph Primitives

Given the generic nature of PBGL, it doesn't contain "primitives" of the sort that one would readily recognize. The lowest level of object that it supports are object types defined within C++ and the STL, and the lowest level objects that can be created within PBGL must have some definitions given in the form of the same datatypes from C++ and STL that it recognizes.

1.6 Execution Model

Figure 1.2: Distributed adjacency list representation



Within PBGL, the primary storage data-structure for graphs is the distributed adjacency list representation. These adjacency lists are partitioned and distributed across multiple computational nodes and the edges and vertices within those partitions are said to "belong to" the computational node upon which its partition of the adjacency list was placed. The same computation node has ownership of the outgoing edges from all of the vertices that it owns as well. Figure 1.2 illustrates the distribution of vertex and edge ownership over three computation nodes.

Graph object properties and property maps are also stored on the computational node local to computational node that owns the graph object. From these data-structures, the executional model is derived from the interplay between the Process Group and the parallel algorithm. Figure 1.3 illustrates the relationship between all of the components and how they interconnect.


Within the process group, PBGL masks MPI message passing with standard BGL flavor methods that maintain data consistency on the backend. Direct commands can be passed to the underlying MPI system for added flexibility, but most graph algorithms will find the process group interface more than adequate.

1.7 Examples

A great deal of a developer's time using the PBGL will be spent developing wrappers to reduce the code complexity of using the libraries directly in all of one's code. For basic examples of usage please see the Section 1.4 on Syntax. This section will illustrate an example of PBGL wrapper class that could be used to lower the cognitive load of using the PBGL.

PBGL

Figure 1.3: Full PBGL Architecture



Listing 1.2: Simple PBGL Wrapper

```

1 #include <boost/config.hpp>
2 #include <boost/version.hpp>
3 #include <boost/graph/graph_utility.hpp>
4 #include <boost/graph/adjacency_list.hpp>
5 #include <boost/property_map/property_map.hpp>
6 #include <boost/static_assert.hpp>
7
8 // Enable PBGL interfaces to BGL algorithms
9 #include <boost/graph/use_mpi.hpp>
10
11 // Communication via MPI
12 #include <boost/graph/distributed/mpi_process_group.hpp>
13
14 // Distributed adjacency list
15 #include <boost/graph/distributed/adjacency_list.hpp>
16
17
18 using namespace boost;
19 using boost::graph::distributed::mpi_process_group;
20 using namespace std;
21
22 // VertexProperties is the property container that will be passed into the
23 // templated graph object - should work for serial and parallel versions
24 struct VertexProperties
25 {
26     std::string name;
27 };
28
29 // EdgeProperties is the property container that will be passed into the
30 // templated graph object - should work for serial and parallel versions
31 struct EdgeProperties
32 {

```

PBGL

```

33     int distance = 0;
34 };
35
36 /* definition of basic boost::graph properties */
37 enum vertex_properties_t { vertex_properties };
38 enum edge_properties_t { edge_properties };
39
40 namespace boost {
41     BOOST_INSTALL_PROPERTY(vertex, properties);
42     BOOST_INSTALL_PROPERTY(edge, properties);
43 }
44
45 // Graph Class
46 template < typename VERTEXPROPERTIES, typename EDGEPROPERTIES >
47 class Graph
48 {
49 public:
50     /* a distributed adjacency list */
51     /* An undirected, weighted graph with distance values stored on the vertices. */
52     typedef adjacency_list<
53         vecS,
54         distributedS<mpi_process_group, vecS>,
55         bidirectionalS, // directed graph
56         property<vertex_properties_t, VERTEXPROPERTIES>,
57         property<edge_properties_t, EDGEPROPERTIES> >
58     GraphContainer;
59
60     /* typedefs to begin normalization of Boost syntax */
61     typedef typename graph_traits<GraphContainer>::vertex_descriptor Vertex;
62     typedef typename graph_traits<GraphContainer>::edge_descriptor Edge;
63     typedef std::pair<Edge, Edge> EdgePair;
64
65     typedef typename graph_traits<GraphContainer>::vertex_iterator vertex_iter;
66     typedef typename graph_traits<GraphContainer>::edge_iterator edge_iter;
67     typedef typename graph_traits<GraphContainer>::adjacency_iterator adjacency_iter;
68     typedef typename graph_traits<GraphContainer>::out_edge_iterator out_edge_iter;
69     typedef typename graph_traits<GraphContainer>::degree_size_type degree_t;
70
71     typedef std::pair<adjacency_iter, adjacency_iter> adjacency_vertex_range_t;
72     typedef std::pair<out_edge_iter, out_edge_iter> out_edge_range_t;
73     typedef std::pair<vertex_iter, vertex_iter> vertex_range_t;
74     typedef std::pair<edge_iter, edge_iter> edge_range_t;
75
76
77     /* constructors etc. */
78     Graph()
79     {}
80
81     Graph(const Graph& g) :
82         graph(g.graph)
83     {}
84
85     virtual ~Graph()
86     {}

```

```

87
88
89  /* structure modification methods */
90  void Clear()
91  {
92      graph.clear();
93  }
94
95  Vertex AddVertex(const VERTEXPROPERTIES& prop)
96  {
97      Vertex v = add_vertex(graph);
98      properties(v) = prop;
99      return v;
100 }
101
102 void RemoveVertex(const Vertex& v)
103 {
104     clear_vertex(v, graph);
105     remove_vertex(v, graph);
106 }
107
108 EdgePair AddEdge(const Vertex& v1, const Vertex& v2,
109                 const EDGEPROPERTIES& prop_12, const EDGEPROPERTIES& prop_21)
110 {
111     /* TODO: maybe one wants to check if this edge could be inserted */
112     Edge addedEdge1 = add_edge(v1, v2, graph).first;
113     Edge addedEdge2 = add_edge(v2, v1, graph).first;
114
115     properties(addedEdge1) = prop_12;
116     properties(addedEdge2) = prop_21;
117
118     return EdgePair(addedEdge1, addedEdge2);
119 }
120
121 Edge AddDirectedEdge(const Vertex& v1,
122                     const Vertex& v2, const EDGEPROPERTIES& prop)
123 {
124     return add_edge(v1, v2, graph).first;
125 }
126
127 /* selectors and properties */
128
129 adjacency_vertex_range_t getAdjacentVertices(const Vertex& v) const
130 {
131     return adjacent_vertices(v, graph);
132 }
133
134 int getVertexCount() const
135 {
136     return num_vertices(graph);
137 }
138
139 int getVertexDegree(const Vertex& v) const
140 {

```



```
141     return degree(v, graph);
142   }
143
144 protected:
145     GraphContainer graph;
146 };
```

1.8 Conclusion

Forthcoming.

Bibliography

- [1] C++ Boost. Libraries, 2012.
- [2] Douglas Gregor and Andrew Lumsdaine. The parallel bgl: A generic library for distributed graph computations. *Parallel Object-Oriented Scientific Computing (POOSC)*, 2:1–18, 2005.