Pragel: A system for Large-Scale Graph Processing Malewicz, Austern, Bik, Dehnert, Horn, Leiser, Czajkowski (Google – 2010)

Sualeh Asif

6.886 Talk MIT Computer Science Artificial Intelligence Laboratory Massachusetts Institute of Technology

March 30, 2021



A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING.



A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING.

GRAPH PROCESSING

PRAGEL : DISTRIBUTED GRAPH PROCESSING A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING. Options: RECAP: (Large-scale) GRAPH PROCESSING

PRAGEL : DISTRIBUTED GRAPH PROCESSING A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING. Options: RECAP: * Shared Memory : LIGRA, Julienne (Large-scale) GRAPH PROCESSING





PRAGEL : DISTRIBUTED GRAPH PROCESSING A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING. Options: RECAP: * Shared Memory: LIGRA, Julienne Boust (BGL), LEDA, Networky, JDSL, FGL con: single computer (Large-scale) GRAPH PROCESSING Parallel / distributed : Parallel BGL Systems : CGM graph con: not good on fault-tolerencel Map Reduce : < no framework > bod usability, bad performance con : ill-suited for most graph Do it yourself....



















Model: * iterate in supersteps RAGEL : DISTRIBUTED GRAPH PROCESSING * operate on all 'utx's (in parallel) in each A SYSTEM FOR LARGE-SCALE GRAPH PROCESSING. MODEL : message possing interface! step Round S : * MORE DETAILS AND (C++) API Every vertex, V, can * read $msg \rightarrow V$ * send $msg \rightarrow W \in ngh(v)$ Halting: * modify state each step, a vtx can vote to halt! Program termintes when * vertices are halted! vtx can be revived if * it recieves a msg. Gevery vtx has * Unique D. * mutable (user-defined) value. edge - list















Model: * iterate in supersteps * operate on all vtx's (in parallel) in each RAGEL : DISTRIBUTED GRAPH PROCESSING FOR LARGE-SCALE GRAPH PROCESSING. A SYSTEM possing interface! MODEL message sted * vertex, V, can Example fiend more in engh(v) Superstep 0 * modify state Vote to halt template <typename VertexValue, typename EdgeValue, typename MessageValue> class Vertex { public: virtual void Compute(MessageIterator* msgs) = 0; const string& vertex_id() const; int64 superstep() const; const VertexValue& GetValue(); VertexValue* MutableValue(); OutEdgeIterator GetOutEdgeIterator(); void SendMessageTo(const string& dest_vertex, const MessageValue& message); void VoteToHalt(); 1:

Model: * iterate in supersteps * operate on all vtx's (in parallel) in each PRAGEL : DISTRIBUTED GRAPH PROCESSING FOR LARGE-SCALE GRAPH PROCESSING. A SYSTEM possing interface! MODEL message sted * vertex, V, can Example Halti Fiend more $\rightarrow w \in ngh(v)$ Superstep * modify state Vote to halt template <typename VertexValue, typename EdgeValue, typename MessageValue> class Vertex { public: virtual void Compute(MessageIterator* msgs) = 0; const string& vertex_id() const; int64 superstep() const; const VertexValue& GetValue(): VertexValue* MutableValue(); OutEdgeIterator GetOutEdgeIterator(); void SendMessageTo(const string& dest_vertex, const MessageValue& message); void VoteToHalt(); 1:



Model: * iterate in supersteps * operate on all vtx's (in parallel) in each PRAGEL : DISTRIBUTED GRAPH PROCESSING FOR LARGE-SCALE GRAPH PROCESSING. A SYSTEM possing interface! MODEL message sted * vertex, V, can Example fiend more in engh(v) Superstep 2 * modify state Vote to halt template <typename VertexValue, typename EdgeValue, typename MessageValue> class Vertex { public: virtual void Compute(MessageIterator* msgs) = 0; const string& vertex_id() const; int64 superstep() const; const VertexValue& GetValue(); VertexValue* MutableValue(); OutEdgeIterator GetOutEdgeIterator(); void SendMessageTo(const string& dest_vertex, const MessageValue& message); void VoteToHalt(); 1:

Model: * iterate in supersteps * operate on all vtx's (in parallel) in each PRAGEL : DISTRIBUTED GRAPH PROCESSING FOR LARGE-SCALE GRAPH PROCESSING. A SYSTEM passing interface! MODEL message sted * vertex, V, can Example fiend more in engh(v) Superstep 2 * modify state Vote to halt template <typename VertexValue, typename EdgeValue, typename MessageValue> class Vertex { public: virtual void Compute(MessageIterator* msgs) = 0; const string& vertex_id() const; int64 superstep() const; const VertexValue& GetValue(); VertexValue* MutableValue(); OutEdgeIterator GetOutEdgeIterator(); void SendMessageTo(const string& dest_vertex, const MessageValue& message); void VoteToHalt(); 1:















Mode : iterate in supersteps operate on all vtx's in parallel) in each PRAGEL : DISTRIBUTED GRAPH PROCESSING ¥ * operate on FOR LARGE-SCALE GRAPH PROCESSING. A SYSTEM (in parallel) in APPLICATIONS ! : step finally * vertex, V, can $> W \in ngh(v)$ modify Vote to halt template <typename VertexValue, typename EdgeValue, typename MessageValue> class Vertex { public: virtual void Compute(MessageIterator* msgs) = 0; const string& vertex_id() const; int64 superstep() const; const VertexValue& GetValue(); VertexValue* MutableValue(); OutEdgeIterator GetOutEdgeIterator(); void SendMessageTo(const string& dest_vertex, const MessageValue& message); void VoteToHalt(); };



```
Mode :
      EL : DISTRIBUTED GRAPH PROCESSING
                                                                                         ¥
                                                                                                             in superstep
                                                                                              iterate
                                                                                         *
                                                                                             operate
                              GE-SCALE GRAPH PROCESSING.
SYSTEM
                                                                                            (in parallel)
                                 >inherits from Vertex
                                                                                               sted
                                                                                         *
   PAGE
                                                                                                       vertex . V. can
                                        Vertex
              class PageRankVertex
                                                                                                                 W \in ngh(V)
                  : public Vertex<double, void, double> {
              public:
                                                                                                 modity
                virtual void Compute(MessageIterator* msgs) {
                  if (superstep() >= 1) {
                    double sum = 0;
                    for (; !msgs->Done(); msgs->Next())
                                                                                                         Vote to halt
                       sum += msgs->Value();
                    *MutableValue() =
                                                                                         template <typename VertexValue,
                         0.15 / NumVertices() + 0.85 * sum;
                                                                                                typename EdgeValue,
                                                                                                typename MessageValue>
                  }
                                                                                         class Vertex {
                                                                                          public:
                                                                                           virtual void Compute(MessageIterator* msgs) = 0;
                  if (superstep() < 30) {
                                                                                           const string& vertex_id() const;
                    const int64 n = GetOutEdgeIterator().size();
                                                                                           int64 superstep() const;
                    SendMessageToAllNeighbors(GetValue() / n);
                                                                                           const VertexValue& GetValue();
                  } else {
                                                                                           VertexValue* MutableValue();
                    VoteToHalt();
                                                                                           OutEdgeIterator GetOutEdgeIterator();
                  }
                                                                                           void SendMessageTo(const string& dest_vertex,
                                                                                                        const MessageValue& message);
                                                                                           void VoteToHalt();
              };
                                                                                         1:
```

```
Mode :
RAGEL : DISTRIBUTED GRAPH PROCESSING
                                                                                        ¥
                                                                                                           in supersteps
                                                                                             iterate.
                                                                                        *
                                                                                           operate
                               GE-SCALE GRAPH PROCESSING.
  SYSTEN
                                                                                          (in parallel) in
                                  >inherits from Vertex
                                                                                             sted
                                                                                       *
     PAGE
                                                                                                     vertex . V. can
                                         vertex
               class PageRankVertex
                                                                                                               W \in udh(v)
                    : public Vertex<double, void. double>
             OHublic:
                                                                                               modity
                 virtual void Compute(MessageIterator* msgs)
                                                                   109
              double sum = 0:
                     for (; !msgs->Done(); msgs->Next())
                                                                                                       Vote to halt
                        sum += msgs->Value();
                      *MutableValue() =
  COMDU
                                                                      erank
                                                                                        template <typename VertexValue,
                          0.15 / NumVertices() + 0.85 * sum;
                                                                                              typename EdgeValue,
                                                                                              typename MessageValue>
                    7
                                                                                        class Vertex {
                                                                                        public:
                                                                                         virtual void Compute(MessageIterator* msgs) = 0;
                   if (superstep() < 30) {
                                                                                         const string& vertex_id() const;
                      const int64 n = GetOutEdgeIterator().size();
                                                                                         int64 superstep() const;
                      SendMessageToAllNeighbors(GetValue() / n);
                                                                                         const VertexValue& GetValue():
                   } else {
                                                                                         VertexValue* MutableValue();
                      VoteToHalt();
                                                                                         OutEdgeIterator GetOutEdgeIterator();
                                                                                         void SendMessageTo(const string& dest_vertex,
                                                                                                      const MessageValue& message);
                                                                                         void VoteToHalt();
               };
                                                                                        1:
```

```
Mode :
PRAGEL : DISTRIBUTED GRAPH PROCESSING
                                                                                                           in supersteps
                                                                                         ¥
                                                                                             iterate
                                                                                        * operate on all vtx's
(in parallel) in each
                              REE-SCALE GRAPH PROCESSING.
A SYSTEM
                                  >inherits from Vertex
                                                                                              step
                                                                                        *
      PAGE
                                                                                                     vertex V, can
                                          vertex
                class PageRankVertex
                                                                                                              >WE nqh(V)
                    : public Vertex<double, void, double> {
              O Mublic:
virtual void Compute(MessageIterator* msgs)
                                                                                             * modity state
                                                                  ~ 109
              double sum = 0:
                      for (; !msgs->Done(); msgs->Next())
                                                                                                        Vote to halt
                        sum += msgs->Value();
                      *MutableValue() =
   COMPU
                                                                       erank
                                                                                         template <typename VertexValue,
                          0.15 / NumVertices() + 0.85 * sum;
                                                                                               typename EdgeValue,
                                                                                               typename MessageValue>
                                                                                         class Vertex {
                                                                                         public:
                                                                                          virtual void Compute(MessageIterator* msgs) = 0;
                    if (superstep() < 30) {
                      const int64 n = GetOutEdgeIterator().size();
                                                                                          const string& vertex_id() const;
                                                                                          int64 superstep() const;
                      SendMessageToAllNeighbors(GetValue() / n);
                                                                                          const VertexValue& GetValue():
                    } else {
    rounds C-
                                                                          message
                                                                                          VertexValue* MutableValue();
                      VoteToHalt();
                                                                                          OutEdgeIterator GetOutEdgeIterator():
                                                                                          void SendMessageTo(const string& dest_vertex,
                                                                                                       const MessageValue& message);
                                                                                          void VoteToHalt();
                };
                                                                                         1:
                                after 30 rounds.
```

```
Mode :
       EL : DISTRIBUTED GRAPH PROCESSING
                                                                                          ¥
                                                                                               iterate
                                                                                                              in superstep
                                                                                          *
                                                                                              operate
             AORD NARGE-SCALE GRAPH PROCESSING.
SYSTEM
                                                                                             (in parallel
                                                                                                sted
                                                                                          *
                                                                                                        vertex, V, can
                                                                                                                   WE no
              class ShortestPathVertex
                                                                                                  modity
                   : public Vertex<int, int, int> {
                void Compute(MessageIterator* msgs) {
                  int mindist = IsSource(vertex_id()) ? 0 : INF;
                  for (; !msgs->Done(); msgs->Next())
                                                                                                          Vote to halt
                     mindist = min(mindist, msgs->Value());
                  if (mindist < GetValue()) {
                                                                                          template <typename VertexValue,
                                                                                                 typename EdgeValue,
                     *MutableValue() = mindist:
                                                                                                 typename MessageValue>
                     OutEdgeIterator iter = GetOutEdgeIterator();
                                                                                          class Vertex {
                                                                                           public:
                    for (; !iter.Done(); iter.Next())
                                                                                            virtual void Compute(MessageIterator* msgs) = 0;
                       SendMessageTo(iter.Target(),
                                                                                            const string& vertex_id() const;
                                      mindist + iter.GetValue());
                                                                                            int64 superstep() const;
                   3
                                                                                            const VertexValue& GetValue();
                  VoteToHalt();
                                                                                            VertexValue* MutableValue();
                                                                                            OutEdgeIterator GetOutEdgeIterator();
                }
              };
                                                                                            void SendMessageTo(const string& dest_vertex,
                                                                                                         const MessageValue& message);
                                                                                            void VoteToHalt();
                                                                                          1:
```

```
Mode :
        GEL : DISTRIBUTED GRAPH PROCESSING
                                                                                                               in supersteps
                                                                                            ¥
                                                                                                 iterate.
                                                                                                operate on
                  FOR LARGE-SCALE GRAPH PROCESSING
                                                                                               (in parallel) in
    SYSTEM
                   ATIONS! .
                                                                                                  step
                                                                                            *
                                                                                                         vertex . V. can
                                            >vtx value : init +INF
>Edge Value : (weighted)
                                                                                            Halt
                                                                                                                  ≫WE nqh(V)
Vertex
                  class ShortestPathVertex
                                                                                                 * modity
                        public Vertex<int, int, int> {
recieves
             Dotentiabid Compute (MessageIterator* msgs) {
                      int mindist = IsSource(vertex_id()) ? 0 : INF;
min distances
                      for (; !msgs->Done(); msgs->Next())
                                                                                                            Vote to halt
                        mindist = min(mindist, msgs->Value());
         Source
from
                      if (mindist < GetValue()) {</pre>
                                                                                            template <typename VertexValue,
                                                                                                   typename EdgeValue,
                         *MutableValue() = mindist:
                                                                                                   typename MessageValue>
                        OutEdgeIterator iter = GetOutEdgeIterator();
                                                                                            class Vertex {
                                                                                             public:
                        for (; !iter.Done(); iter.Next())
                                                                                              virtual void Compute(MessageIterator* msgs) = 0;
                           SendMessageTo(iter.Target(),
                                                                                              const string& vertex_id() const;
                                         mindist + iter.GetValue());
                                                                                              int64 superstep() const;
                                                                                              const VertexValue& GetValue();
                      VoteToHalt();
                                                                                              VertexValue* MutableValue();
                                                                                              OutEdgeIterator GetOutEdgeIterator():
                    3
                  };
                                                                                              void SendMessageTo(const string& dest_vertex,
                                                                                                          const MessageValue& message);
                                                                                              void VoteToHalt();
                                                                                            1:
```



Figure 7: SSSP—1 billion vertex binary tree: varying number of worker tasks scheduled on 300 multicore machines



Figure 8: SSSP—binary trees: varying graph sizes on 800 worker tasks scheduled on 300 multicore machines



