Graph matching for BSP

Rob H. Bisseling

Mathematical Institute, Utrecht University
Joint work with Fredrik Manne (Bergen, Norway)

KU Leuven, October 25, 2013

Matching

- Introduction
- Greedy matching

BSP algorithm for edge-weighted matching

- Sequential approximation algorithm
- BSP approximation algorithm

Conclusion

Matchmaker, Matchmaker, Make me a match

From the film *Fiddler on the roof*

- Hodel: Well, somebody has to arrange the matches. Young people can't decide these things themselves.
- Hodel: For Papa, make him a scholar.
- Chava: For Mama, make him rich as a king.

Matching can win you a Nobel prize

Marriage as an Economic Problem

Lloyd Shapley and Alvin Roth win the Nobel Prize for showing the best way to match people with what they really want.

By Matthew Yglesias | Posted Monday, Oct. 15, 2012, at 1:51 PM ET

65 SOCIAL OFF Like 963 Tweet 100

The Nobel Prize in economics went to Alvin E. Roth and Lloyd S. Shapley "for the theory of stable allocations and the practice of market design"

Source: Slate magazine October 15, 2012

Universiteit Utrecht

Motivation of graph matching

- Graph matching is a pairing of neighbouring vertices within a graph.
- Matching has applications in
  - medicine: finding suitable donors for organs
  - social networks: finding friends or partners
  - scientific computing: finding pivot elements in matrix computations
  - graph coarsening in multilevel methods:
    make the problem smaller by merging similar vertices before partitioning it for parallel computing
  - Finding similarity in Protein-Protein Interaction networks

Motivation of greedy/approximation graph matching

- Optimal solution is possible in polynomial time: not NP-hard.
- Time for weighted matching in graph $G = (V, E)$ is $\mathcal{O}(mn + n^2 \log n)$ with $n = |V|$ the number of vertices, and $m = |E|$ the number of edges (Gabow 1990).
- The aim is $n = 10^9$, and perhaps $m = 10^{11}$, so a time of $\mathcal{O}(10^{20})$ is far too long. We need linear-time greedy algorithms, or approximation algorithms.

Formal definition of graph matching

- A graph is a pair $G = (V, E)$ with vertices V and edges E.
- All edges $e \in E$ are of the form $e = (v, w)$ for vertices $v, w \in V$.
- A matching is a collection $M \subseteq E$ of disjoint edges.
- Here, the graph is undirected, so $(v, w) = (w, v)$.

Maximal matching

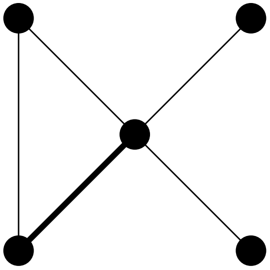

- A matching is maximal if we cannot enlarge it further by adding another edge to it.

Maximum matching

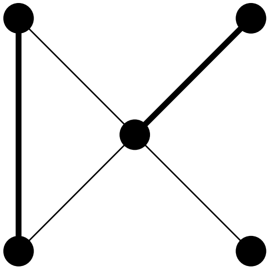

- A matching is maximum if it possesses the largest possible number of edges, compared to all other matchings.

Edge-weighted matching

- If the edges are provided with weights $\omega : E \to \mathbb{R}_{>0}$, finding a matching M which maximises

$$\omega(M) = \sum_{e \in M} \omega(e),$$

  is called edge-weighted matching.
- Greedy matching provides us with maximal matchings, but not necessarily with maximum possible weight.

Sequential greedy matching

- We will first look at a sequential greedy algorithm which generates a maximal matching.
- In random order, vertices $v \in V$ select and match neighbours one-by-one.
- Here, we can pick
  - the first available neighbour w of v (greedy random matching)
  - the neighbour w with maximum $\omega(v, w)$ (greedy weighted matching)
- Or: we sort the edges by weight, and successively match the vertices v and w of the heaviest available edge (v, w) (greedy matching)

Sequential greedy matching

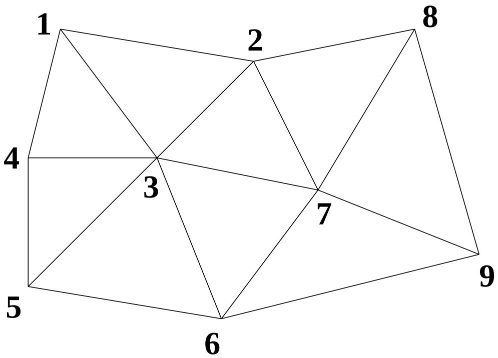

Universiteit Utrecht

Sequential greedy matching

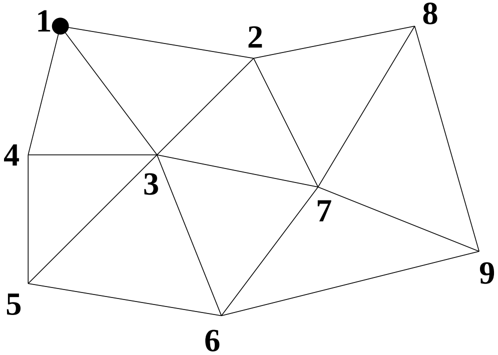

Universiteit Utrecht

Sequential greedy matching

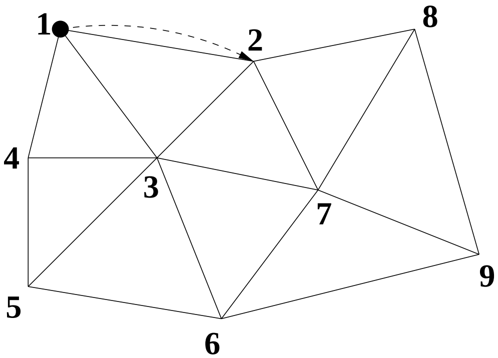

Sequential greedy matching

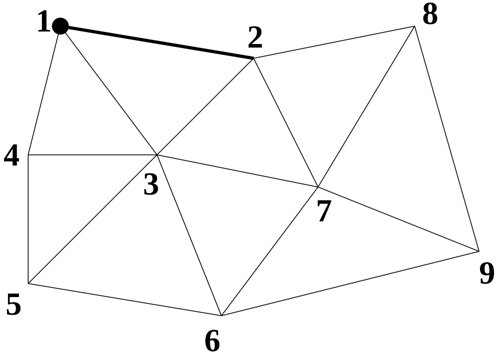

Sequential greedy matching

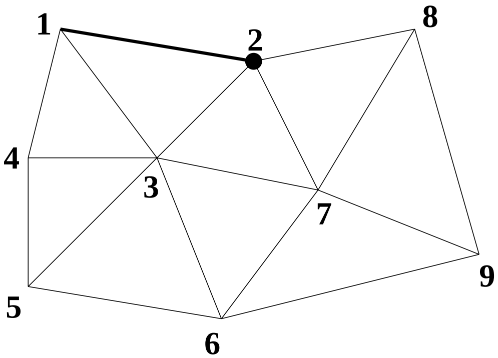

Sequential greedy matching

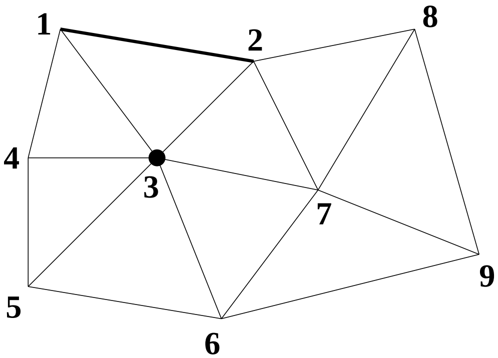

Universiteit Utrecht

Sequential greedy matching

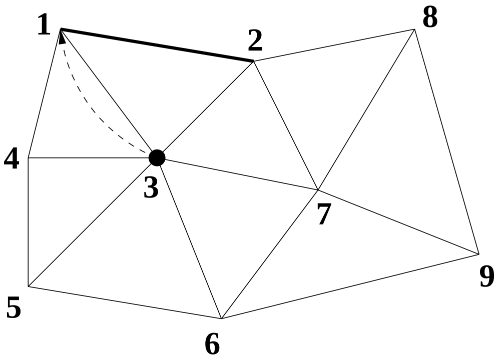

Sequential greedy matching

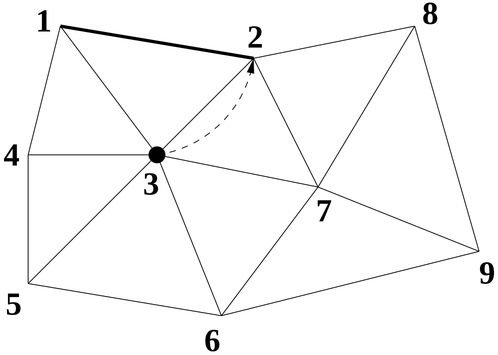

Sequential greedy matching

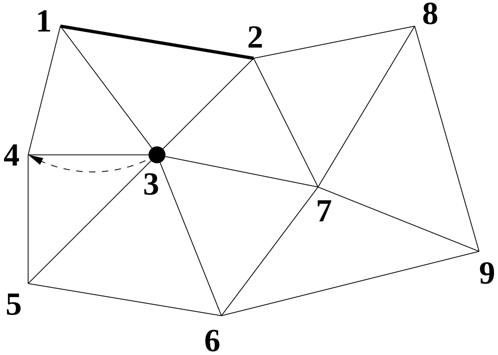

Sequential greedy matching

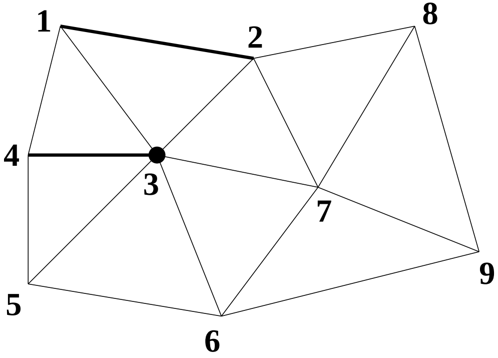

Universiteit Utrecht

Sequential greedy matching

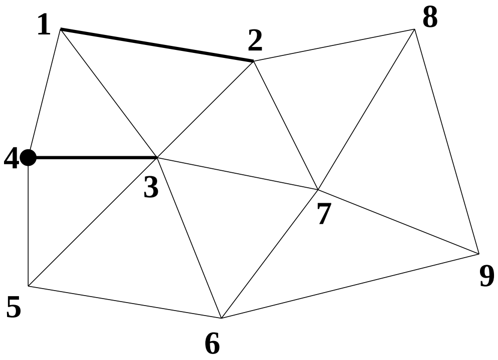

Universiteit Utrecht

Sequential greedy matching

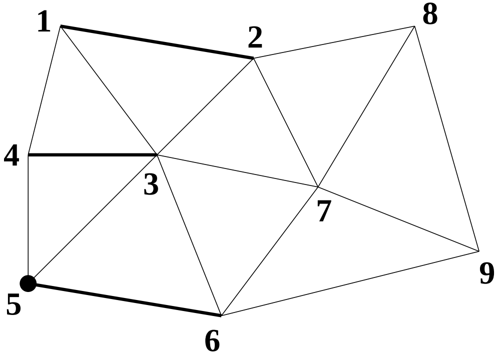

Sequential greedy matching

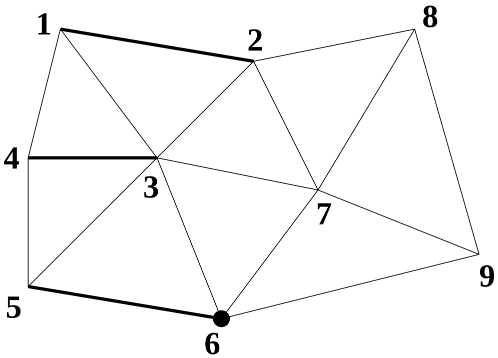

Sequential greedy matching

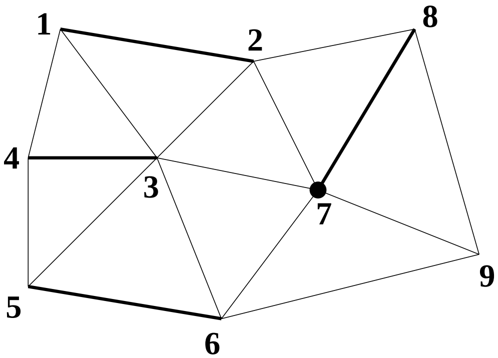

Sequential greedy matching

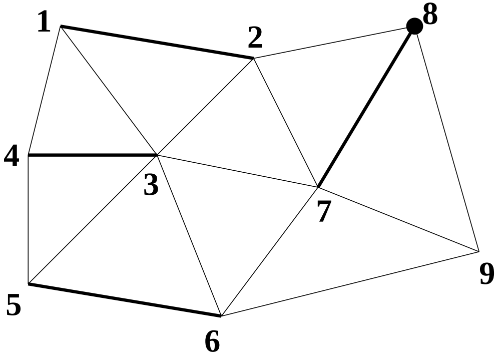

Sequential greedy matching

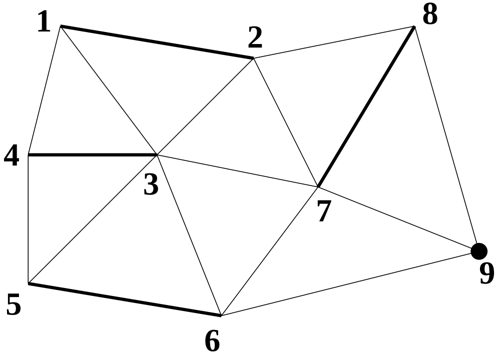

Sequential greedy matching

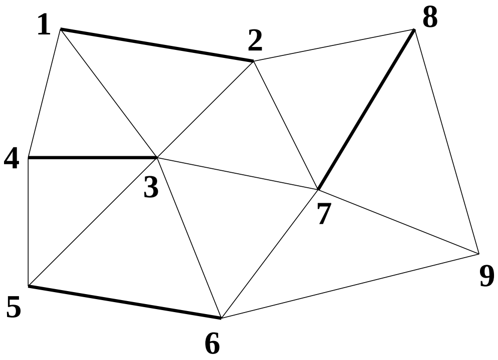

Greedy is a 1/2-approximation algorithm

- Weight $\omega(M) \geq \omega_{\text{optimal}}/2$
- Cardinality $|M| \geq |M_{\text{card}-\text{max}}|/2$, because M is maximal.

Parallel greedy matching: trouble

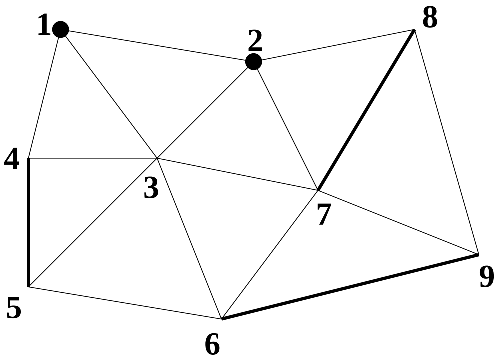

Suppose we match vertices simultaneously.

Parallel greedy matching: trouble

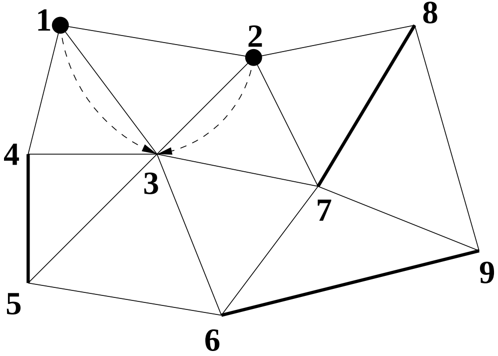

Two vertices each find an unmatched neighbour. . .

Parallel greedy matching: trouble

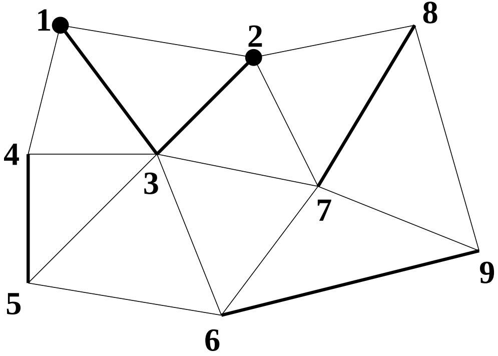

. . . but generate an invalid matching.

Dominant-edge algorithm

while $E \neq \emptyset$ **do**
  pick dominant edge $(v, w) \in E$
  $M := M \cup \{(v, w)\}$
  $E := E \setminus \{(x, y) \in E : x = v \vee x = w\}$
  $V := V \setminus \{v, w\}$
return M

- An edge $(v, w) \in E$ is dominant if

$$\omega(v, w) = \max\{\omega(x, y) : (x, y) \in E \wedge (x = v \vee x = w)\}$$

Proof: algorithm is 1/2-approximation

- Let M be the matching produced by the dominant-edge algorithm.
- Let M^* be a maximum matching with weight ω_{optimal}.
- Let $e^* = (a, b) \in M^*$. If both a and b are not matched in M, we could add (a, b) to M. Hence M is not maximal.
- Contradiction, because the algorithm only terminates when it produces a maximal matching.
- Therefore, a or b is matched in M.

Proof (cont'd)

- Let $M^* = \{e_0^*, \ldots, e_{k-1}^*\}$. For each edge $e_i^* \in M^*$, define an edge $e_i \in M$, as follows. If $e_i^* \in M$, $e_i = e_i^*$, otherwise e_i is the edge that removes e_i^* from E in the algorithm.
- It may happen that $e_i = e_j$ for $i \neq j$.
- $\omega(e_i) \geq \omega(e_i^*)$ for all i, since e_i is locally dominant in the algorithm and removes e_i^*, or $e_i = e_i^*$.
- Every edge $e \in M$ can occur at most twice in the list of e_i's, since it can remove from E at most 2 edges from M^*.
- $$2\omega(M) \geq \sum_{i=0}^{k-1} \omega(e_i) \geq \sum_{i=0}^{k-1} \omega(e_i^*) = \omega_{\mathrm{optimal}}$$
- Hence $\omega(M) \geq \omega_{\mathrm{optimal}}/2$.

Sequential approximation algorithm: initialisation

function $\textsc{SeqMatching}(V, E)$
 for all $v \in V$ **do**
  $pref(v) = null$
 $D := \emptyset$
 $M := \emptyset$

 { Find dominant edges }
 for all $v \in V$ **do**
  $Adj_v := \{w \in V : (v, w) \in E\}$
  $pref(v) := \mathrm{argmax}\{\omega(v, w) \; : \; w \in Adj_v\}$
  if $pref(pref(v)) = v$ **then**
   $D := D \cup \{v, pref(v)\}$
   $M := M \cup \{(v, pref(v))\}$

Sequential approximation algorithm: main loop

while $D \neq \emptyset$ **do**
  pick $v \in D$
  $D := D \setminus \{v\}$
  for all $x \in Adj_v \setminus \{pref(v)\} \ : \ (x, pref(x)) \notin M$ **do**
    $Adj_x := Adj_x \setminus \{v\}$
    $pref(x) := \mathrm{argmax}\{\omega(x, w) \ : \ w \in Adj_x\}$
    if $pref(pref(x)) = x$ **then**
      $D := D \cup \{x, pref(x)\}$
      $M := M \cup \{(x, pref(x))\}$
return M

Dominant-edge algorithm

- Based on dominant edges:
  $v = pref(w)$ and $w = pref(v)$.
- Dominance is a local property: easy to parallelise.
- Algorithm keeps going until set of dominant vertices D is empty and matching M is maximal.
- Assumption without loss of generality: weights are unique.
  Otherwise, add $r \cdot \epsilon$ to weight, with $r \in [0, 1]$ random, or use vertex numbering to break ties.

Time complexity

- Linear time complexity $\mathcal{O}(|E|)$ if edges of each vertex are sorted by weight.
- Sorting costs are

$$\sum_v deg(v) \log deg(v) \leq \sum_v deg(v) \log \Delta = 2|E| \log \Delta,$$

  where Δ is the maximum vertex degree.
- This algorithm is based on a dominant-edge algorithm by Preis (1999), called LAM, which is linear-time $\mathcal{O}(|E|)$, does not need sorting, and also is a 1/2-approximation, but is hard to parallelise.

Parallel algorithm (Manne & Bisseling, 2007)

- Processor $P(s)$ has vertex set V_s, with

$$\bigcup_{s=0}^{p-1} V_s = V$$

and $V_s \cap V_t = \emptyset$ if $s \neq t$.

- The adjacency set Adj_v of a vertex v may contain vertices w from another processor.

Halo vertices

- We define the set of halo vertices

$$H_s = \bigcup_{v \in V_s} Adj_v \setminus V_s$$

- The weights $\omega(v, w)$ are stored with the edges, for all $v \in V_s$ and $w \in V_s \cup H_s$.
- $E_s = \{(v, w) \in E : v \in V_s\}$
  is the subset of all the edges connected to V_s.

Parallel algorithm for $P(s)$: initialisation

function $\textsc{ParMatching}(V_s, H_s, E_s, \text{distribution } \phi)$
  for all $v \in V_s$ **do**
    $pref(v) = null$
  $D_s := \emptyset$
  $M_s := \emptyset$

  { Find dominant edges }
  for all $v \in V_s$ **do**
    $Adj_v := \{w \in V_s \cup H_s : (v, w) \in E_s\}$
    $pref(v) := \operatorname{argmax}\{\omega(v, w) \ : \ w \in Adj_v\}$
    if $pref(v) \in V_s$ **then**
      if $pref(pref(v)) = v$ **then**
        $D_s := D_s \cup \{v, pref(v)\}$
        $M_s := M_s \cup \{(v, pref(v))\}$
    else
      put $proposal(v, pref(v))$ in $P(\phi(pref(v)))$
  Sync

How to propose

Source: www.theguardian.com

proposal(v, w): v proposes to w

Parallel algorithm for $P(s)$: main loop

while $D_s \neq \emptyset$ **do**
  pick $v \in D_s$
  $D := D \setminus \{v\}$
  for all $x \in Adj_v \setminus \{pref(v)\} \ : \ (x, pref(x)) \notin M_s$ **do**
    if $x \in V_s$ **then**
      $Adj_x := Adj_x \setminus \{v\}$
      $pref(x) := \text{argmax}\{\omega(x, w) \ : \ w \in Adj_x\}$
      if $pref(x) \in V_s$ **then**
        if $pref(pref(x)) = x$ **then**
          $D_s := D_s \cup \{x, pref(x)\}$
          $M_s := M_s \cup \{(x, pref(x))\}$
      else
        put $proposal(x, pref(x))$ in $P(\phi(pref(x)))$
    else $\{x \in H_s\}$
      put $unavailable(v, x)$ in $P(\phi(x))$
Sync

Universiteit Utrecht

Parallel algorithm for $P(s)$: communication (1)

for all messages m received **do**
  if $m = proposal(x, y)$ **then**
    if $pref(y) = x$ **then**
      $D_s := D_s \cup \{y\}$
      $M_s := M_s \cup \{(x, y)\}$
      put $accepted(x, y)$ in $P(\phi(x))$
  if $m = accepted(x, y)$ **then**
    $D_s := D_s \cup \{x\}$
    $M_s := M_s \cup \{(x, y)\}$
  ...

Parallel algorithm for $P(s)$: communication (2)

for all messages m received **do**
  . . .
  if $m = unavailable(v, x)$ **then**
    $Adj_x := Adj_x \setminus \{v\}$
    $pref(x) := \text{argmax}\{\omega(x, w) \ : \ w \in Adj_x\}$
    if $pref(x) \in V_s$ **then**
      if $pref(pref(x)) = x$ **then**
        $D_s := D_s \cup \{x, pref(x)\}$
        $M_s := M_s \cup \{(x, pref(x))\}$
    else
      put $proposal(x, pref(x))$ in $P(\phi(pref(x)))$
Sync

Termination

- The algorithm alternates supersteps of computation running the main loop and communication handling the received messages.
- The whole algorithm can terminate when no messages have been received by processor $P(s)$ and the local set D_s is empty, for all s.
- This can be checked at the every synchronisation point.

Load balance

- Processors can have different amounts of work, even if they have the same number of vertices or edges.
- Use can be made of a global clock based on ticks, the unit of time needed to handle a vertex x (in $\mathcal{O}(1)$).
- After every k ticks, everybody synchronises.
- The BSP system takes care that messages are automatically sent, in bulk.
- In the next superstep, all received messages are handled.

Synchronisation frequency

- Guidance for the choice of k is provided by the BSP parameter l: choosing $k \geq l$ guarantees that at most 50% of the total time is spent in synchronisation.
- Choosing k sufficiently small will cause all processors to be busy during most supersteps.
- Good choice: $k = 2l$?

Further improvement: edge-based (2D) distribution

Name	SpMV 1D	SpMV 2D	Matching 1D	Matching 2D
rw9 (af_shell10)	113	105	169	150
rw10 (boneS10)	150	145	228	189
rw11 (Stanford)	340	141	479	234
rw12 (gupta3)	710	44	1,305	61
rw13 (St_Berk.)	716	448	1,152	812
rw14 (F1)	139	130	148	139
sw1 (small world)	1,007	417	2,111	303
sw2	1,957	829	3,999	563
sw3	2,017	832	4,255	528
er1 (random)	1,856	1,133	1,788	1,157
er2	3,451	1,841	3,721	1,635
er3	5,476	2,569	6,350	1,990

Communication volume in sparse matrix–vector multiplication and Karp–Sipser matching.

Source: Patwary, Bisseling, Manne (2010).

Conclusions and outlook

- BSP is extremely suitable for parallel graph computations:
  - no need to worry about communication because we buffer messages until the next synchronisation;
  - no need for send-receive pairs;
  - BSP cost model gives synchronisation frequency;
  - correctness proof of algorithm becomes simpler;
  - no deadlock possible.