Big graphs for big data: parallel matching and clustering on billion-vertex graphs

Rob H. Bisseling

Mathematical Institute, Utrecht University

Collaborators: Bas Fagginger Auer, Fredrik Manne, Albert-Jan Yzelman

Asia-trip A-Eskwadraat, July 2014



Matching Introduction

Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



Graph Matching

Introduction Greedy algorithm Parallelisable 1/2-approximation algorithm BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential algorithm GPU algorithm Results

Conclusion

Outline

Matching Introduction Greedy Parallelisable BSP algorithm

Clustering Introduction Sequential Results



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



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

・ロト ・ 同ト ・ ヨト ・ ヨ

Universiteit Utrecht

Outlin

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Motivation of graph matching

- Graph matching is a pairing of neighbouring vertices.
- It has applications in
 - medicine: finding suitable donors for organs
 - social networks: finding partners
 - scientific computing: finding pivot elements in matrix computations
 - graph coarsening: making the graph smaller by merging similar vertices

Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential Results



Motivation of greedy/approximation graph matching

- Optimal solution is possible in polynomial time.
- ► Time for weighted matching in graph G = (V, E) is O(mn + n² log n) with n = |V| the number of vertices, and m = |E| the number of edges (Gabow 1990).
- ► The aim is a billion vertices, n = 10⁹, with 100 edges per vertex, i.e. m = 10¹¹.
- ► Thus, a time of O(10²⁰) = 100,000 Petaflop units is far too long. Fastest supercomputer today, the Chinese Tianhe-2 (Milky-Way 2), performs 33.8 Petaflop/s.
- We need linear-time greedy or approximation algorithms.

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential Results

Formal definition of graph matching

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

Outline

Matching

Introduction

- Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Maximal matching



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

Outlin

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results



Maximum matching



Outlin

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

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 ω : E → 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.

Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

Sequential greedy matching

- ► In random order, vertices v ∈ 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 all the edges by weight, and successively match the vertices v and w of the heaviest available edge (v, w). This is commonly called greedy matching.

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion







Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion







Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion







Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion







Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion





Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



Greedy matching is a 1/2-approximation algorithm

- Weight $\omega(M) \ge \omega_{\text{optimal}}/2$
- Cardinality $|M| \ge |M_{\text{card}-\text{max}}|/2$, because M is maximal.
- ► Time complexity is O(m log m), because all edges must be sorted.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Parallel greedy matching: trouble



Suppose we match vertices simultaneously.

- æ

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □



Greedy Parallelicabl

BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

Parallel greedy matching: trouble



Two vertices each find an unmatched neighbour...



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

Parallel greedy matching: trouble



... but generate an invalid matching.

æ

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

13

Parallelisable dominant-edge algorithm

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

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

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





Universiteit Utrecht

Introduction Greedy Parallelisable BSP algorith GPU algorith

Clustering Introduction Sequential Results

Sequential approximation algorithm: initialisation

function 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 $Adj_v := \{w \in V : (v, w) \in E\}$ $pref(v) := \operatorname{argmax}\{\omega(v, w) : w \in Adj_v\}$ if pref(pref(v)) = v then $D := D \cup \{v, pref(v)\}$ $M := M \cup \{(v, pref(v))\}$

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results



Mutual preferences



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



Non-mutual preferences



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



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
 $Adj_x := Adj_x \setminus \{v\}$
 $\{ \text{ Set new preference } \}$
 $pref(x) := \operatorname{argmax}\{\omega(x, w) : w \in Adj_x\}$
if $pref(pref(x)) = x$ then
 $D := D \cup \{x, pref(x)\}$
 $M := M \cup \{(x, pref(x))\}$
return M

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion


Properties of the dominant-edge algorithm

► Dominant-edge algorithm is a 1/2-approximation:

 $\omega(M) \geq \omega_{\rm optimal}/2$

- 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, use vertex numbering to break ties.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Time complexity

- ► Linear time complexity O(|E|) if edges of each vertex are sorted by weight.
- Sorting costs are

$$\sum_{v} deg(v) \log deg(v) \le \sum_{v} deg(v) \log \Delta = 2|E| \log \Delta,$$

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 O(|E|), does not need sorting, and also is a 1/2-approximation, but is hard to parallelise.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Parallel computer: abstract model



Bulk synchronous parallel (BSP) computer. Proposed by Leslie Valiant, 1989.



Universiteit Utrecht

BSP algorithm

Sequential

Parallel algorithm: supersteps



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion



Universiteit Utrecht

Composition with Red, Yellow, Blue and Black



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion

Piet Mondriaan 1921



Universiteit Utrecht

Mondriaan data distribution for matrix prime60



- Non-Cartesian block distribution of 60 × 60 matrix prime60 with 462 nonzeros, for p = 4
- ► $a_{ij} \neq 0 \iff i|j \text{ or } j|i$ (1 ≤ i, j ≤ 60)

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □



Universiteit Utrecht

BSP algorithm

Sequential

Parallel algorithm (Manne & Bisseling, 2007)

• Processor P(s) has vertex set V_s , with

$$igcup_{s=0}^{
ho-1}V_s=V$$

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

This is a p-way partitioning of the vertex set.



- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Halo vertices

- The adjacency set Adj_v of a vertex v may contain vertices w from another processor.
- 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) ∈ E : v ∈ V_s} is the subset of all the edges connected to V_s.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion



Universiteit Utrecht

Parallel algorithm for P(s): initialisation

function PARMATCHING(V_s, H_s, E_s , distribution ϕ) for all $v \in V_s$ do pref(v) = null $D_{s} := \emptyset$ $M_{\mathfrak{s}} := \emptyset$ { Find dominant edges } for all $v \in V_s$ do $Adj_{v} := \{ w \in V_s \cup H_s : (v, w) \in E_s \}$ SetNewPreference(v, Adj_v , pref, V_s , D_s , M_s , ϕ) Sync

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results



Setting a vertex preference

function SETNEWPREFERENCE(v, Adj, V, D, M, ϕ) $pref(v) := \operatorname{argmax}\{\omega(v, w) : w \in Adj\}$ if $pref(v) \in V$ then if pref(pref(v)) = v then $D := D \cup \{v, pref(v)\}$ $M := M \cup \{(v, pref(v))\}$

else

put proposal(v, pref(v)) in $P(\phi(pref(v)))$



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results



How to propose



Source: www.theguardian.com

proposal(v, w): v proposes to w



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



Universiteit Utrecht

Parallel algorithm for P(s): main loop

process received messages

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



- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential



Parallel algorithm for P(s): process received messages

for all messages *m* received do if m = proposal(x, y) then if pref(y) = x then $D_{s} := D_{s} \cup \{y\}$ $M_{\mathfrak{s}} := M_{\mathfrak{s}} \cup \{(x, y)\}$ put accepted(x, y) in $P(\phi(x))$ if m = accepted(x, y) then $D_{s} := D_{s} \cup \{x\}$ $M_{\mathfrak{s}} := M_{\mathfrak{s}} \cup \{(x, y)\}$ if m = unavailable(v, x) then if $(x, pref(x)) \notin M_s$ then $Adi_{\times} := Adi_{\times} \setminus \{v\}$ SetNewPreference(x, Adj_x, pref, V_s , D_s , M_s , ϕ)



Universiteit Utrecht

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Termination

- The algorithm alternates supersteps of computation running the main loop and communication handling the received messages.
- ► The whole algorithm terminates 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 every synchronisation point.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering
- Sequential Results



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 O(1)).
- ► Here, 'handling' could mean setting a new preference.
- After every k ticks, everybody synchronises.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



Synchronisation frequency

- Guidance for the choice of k is provided by the BSP parameter l, the cost of a global synchronisation.
- ► Choosing k ≥ 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 = 2I?



- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion



Universiteit Utrecht

MulticoreBSP enables shared-memory BSP

| Multico | DTEBSP efficient, easy & correct parallelisation | |
|---------------------------|---|--|
| | HOME SOFTWARE COMMUNITY FORUM | Outline |
| Introduction BSP model | Home + Introduction Print this page | Matching |
| BSP library | ECENT NEWS Venion 1.2, and a madmap for future releases. MulticovBDP for 0, venion 1.1 released. see all news items | Introduction Greedy Parallelisable BSP algorithm GPU algorithm |
| | Introduction MulticonsBSP brings Bulk Synchronous Parallel (BSP) programming to modern multicore processors. BSP programming leads to high- performance codes: Speed of SpMV multiplication on a 64-core machine | Clustering Introduction Sequential Results |
| | 10 9 9 9 10 10 10 10 10 10 10 10 10 10 | Conclusion |



- A different approach, tightly coupled to the GPU architecture.
- To prevent matching conflicts, we create two groups of vertices:
 - Blue vertices propose.
 - Red vertices respond.
- Proposals that were responded to, are matched.



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion

Universiteit Utrecht



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential



Greedv GPU algorithm

Sequential





Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



- Greedv GPU algorithm
- Sequential



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential



- Greedv GPU algorithm
- Sequential

Quality of the matching



Outlin

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

Fraction of matched vertices as a function of the number of iterations. Number of edges between 2 and 47 million.



Universiteit Utrecht

Random colouring of the vertices

- At each iteration, we colour the vertices $v \in V$ differently.
- For a fixed $p \in [0, 1]$

$$\mathbf{colour}(v) = \begin{cases} \mathbf{blue} & \text{with probability } p, \\ \mathbf{red} & \text{with probability } 1 - p. \end{cases}$$

- How to choose p? Maximise the number of matched vertices.
- For large random graphs, the expected fraction of matched vertices can be approximated by

$$2(1-p)(1-e^{-\frac{p}{1-p}}).$$

This is independent of the edge density.

Universiteit Utrecht

Outlin

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential

Clustering of road network of the Netherlands





Graph with 2,216,688 vertices and 2,441,238 edges yields 506 clusters with modularity 0.995.

Universiteit Utrecht

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential Results

Formal definition of a clustering

► A clustering of an undirected graph G = (V, E) is a collection C of disjoint subsets of V satisfying

$$V=\bigcup_{C\in\mathcal{C}}C.$$

- Elements $C \in C$ are called clusters.
- The number of clusters is not fixed beforehand.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential Results



Quality measure for clustering: modularity

- The quality measure modularity was introduced by Newman and Girvan in 2004 for finding communities.
- Let G = (V, E, ω) be a weighted undirected graph without self-edges. We define

$$\zeta(\mathbf{v}) = \sum_{(u,v)\in E} \omega(u,v), \qquad \Omega = \sum_{e\in E} \omega(e).$$

• Then, the modularity of a clustering C of G is defined by

$$\mathsf{mod}(\mathcal{C}) = \frac{\sum_{C \in \mathcal{C}} \sum_{\substack{(u,v) \in E \\ u,v \in C}} \omega(u,v)}{\Omega} - \frac{\sum_{C \in \mathcal{C}} \left(\sum_{v \in C} \zeta(v)\right)^2}{4\Omega^2}$$

•
$$-\frac{1}{2} \leq \operatorname{mod}(\mathcal{C}) \leq 1.$$

Universiteit Utrecht

Introduction Sequential
Merging clusters: change in modularity

▶ The set of all cut edges between clusters C and C' is

$$cut(C, C') = \{\{u, v\} \in E \mid u \in C, v \in C'\}$$

If we merge clusters C and C' from C into one cluster C ∪ C', then we get a new clustering C' with

$$\operatorname{mod}(\mathcal{C}') = \operatorname{mod}(\mathcal{C}) + \frac{1}{4\Omega^2} \left(4\Omega\omega(\operatorname{cut}(\mathcal{C},\mathcal{C}')) - 2\zeta(\mathcal{C})\zeta(\mathcal{C}') \right),$$

Outime

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion



Universiteit Utrecht

 $\zeta(\mathcal{C}\cup\mathcal{C}')=\zeta(\mathcal{C})+\zeta(\mathcal{C}').$

Agglomerative greedy clustering heuristic

$$\begin{array}{l} \max \leftarrow -\infty \\ G^{0} = (V^{0}, E^{0}, \omega^{0}, \zeta^{0}) \\ i \leftarrow 0 \\ \mathcal{C}^{0} \leftarrow \{\{v\} \mid v \in V\} \\ \text{while } |V^{i}| > 1 \text{ do} \\ \text{ if } \operatorname{mod}(G, \mathcal{C}^{i}) \geq \max \text{ then} \\ \max \leftarrow \operatorname{mod}(G, \mathcal{C}^{i}) \\ \mathcal{C}^{\operatorname{best}} \leftarrow \mathcal{C}^{i} \\ \mu \leftarrow \operatorname{weighted_match_clusters}(G^{i}) \\ (\pi^{i}, G^{i+1}) \leftarrow \operatorname{coarsen}(G^{i}, \mu) \\ \mathcal{C}^{i+1} \leftarrow \{\{v \in V \mid (\pi^{i} \circ \cdots \circ \pi^{0})(v) = u\} \mid u \in V^{i+1}\} \\ i \leftarrow i + 1 \\ \text{return } \mathcal{C}^{\operatorname{best}} \end{array}$$



Greedv

Sequential

Results: clustering time for DIMACS graphs



Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm
- Clustering Introduction Sequential Results

Conclusion

- DIMACS categories: clustering/, coauthor/, streets/, random/, delaunay/, matrix/, walshaw/, dyn-frames/, and redistrict/.
- CUDA implementation with the Thrust template library and Intel TBB implementation.
- Web link graph uk-2002 with 0.26 billion vertices clustered in 30 s using Intel TBB.



Universiteit Utrecht

DIMACS road networks and coauthor graphs

| G | V | <i>E</i> | mod | t | mod | t | _ |
|---------------|------------|------------|------|------|------|---------------------|---------------------------------|
| | | | CU | CU | TBB | TBB | Outline |
| luxembourg | 114,599 | 119,666 | 0.99 | 0.13 | 0.99 | 0.14 | Matching |
| belgium | 1,441,295 | 1,549,970 | 0.99 | 0.44 | 0.99 | 1.11 | Introduction Greedy |
| netherlands | 2,216,688 | 2,441,238 | 0.99 | 0.62 | 0.99 | 1.72 | Parallelisable BSP algorithm |
| italy | 6,686,493 | 7,013,978 | 1.00 | 1.54 | 1.00 | 5.2 <mark>6</mark> | GPU algorithm |
| great-britain | 7,733,822 | 8,156,517 | 1.00 | 1.79 | 1.00 | 6.0 <mark>0</mark> | Clustering |
| germany | 11,548,845 | 12,369,181 | 1.00 | 2.82 | 1.00 | 9.5 <mark>7</mark> | Sequential |
| asia | 11,950,757 | 12,711,603 | 1.00 | 2.69 | 1.00 | 9.3 <mark>3</mark> | Results |
| europe | 50,912,018 | 54,054,660 | - | | 1.00 | 45. <mark>21</mark> | Conclusion |
| coAuthorsCite | 227,320 | 814,134 | 0.84 | 0.42 | 0.85 | 0. <mark>23</mark> | _ |
| coAuthorsDBLP | 299,067 | 977,676 | 0.75 | 0.59 | 0.76 | 0 <mark>.28</mark> | |
| citationCite | 268,495 | 1,156,647 | 0.64 | 0.89 | 0.68 | 0 <mark>.32</mark> | |
| coPapersDBLP | 540,486 | 15,245,729 | 0.64 | 6.43 | 0.67 | 2.28 | |
| coPapersCite | 434,102 | 16,036,720 | 0.75 | 6.49 | 0.77 | 2.27 | _ |

mod = modularity, t = time in s, CU = CUDA



Conclusions

▶ BSP is extremely suitable for parallel graph computations:

- no worries about communication because we buffer messages until the next synchronisation;
- no send-receive pairs, but one-sided put or get operations;
- BSP cost model gives synchronisation frequency;
- · correctness proof of algorithm becomes simpler;
- no deadlock possible.
- Matching can be the basis for clustering, as demonstrated for GPUs and multicore CPUs.
- ▶ We clustered Asia's road network with 12M vertices and 12.7M edges in 2.7 seconds on a GPU.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering Introduction Sequential Results

Conclusion

