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, Mostofa Patwary, Daan Pelt, Albert-Jan Yzelman

Workshop AMLaGAP, Orléans, May 19, 2014

Outlin

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Graph Matching

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

Clustering

- Introduction Sequential algorithm GPU algorithm Results
- 2D sparse matrix partitioning
- 2D (edge-based) matching
- Conclusion

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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.

Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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

• • • • • • • • • • • • •

Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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 before partitioning it for parallel computations
 - bioinformatics: finding similarity in Protein-Protein Interaction networks

Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



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, $n = 10^9$, with 100 edges per vertex, i.e. $m = 10^{11}$.
- ► Thus, a time of O(10²⁰) = 100,000 Petaflop units is far too long. Fastest supercomputer today, the Tianhe-2, performs 33.8 Petaflop/s.
- We need linear-time greedy or approximation algorithms.



Outline

Matching

Introduction

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

- 2D partitioning
- 2D matching
- Conclusion

References



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 Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



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 Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Maximum matching



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



Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Edge-weighted matching

(日) (部) (注) (注) (三)

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 Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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.

Outline

Matching Introduction Greedy Parallelisable BSP algorithm

GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Uniteduction

Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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 Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Parallel greedy matching: trouble



Suppose we match vertices simultaneously.



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Parallel greedy matching: trouble



Two vertices each find an unmatched neighbour...

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Parallel greedy matching: trouble



... but generate an invalid matching.

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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

Greedy Parallelisable

Sequential GPU algori 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 Results

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Mutual preferences



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Non-mutual preferences



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References


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



Universiteit Utrecht

Greedy Parallelisable

Clustering Introduction Sequential GPU algorit Results

Properties of the dominant-edge algorithm

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

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

Dominant edge means mutual preference:

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, use vertex numbering to break ties.

Universiteit Utrecht

Parallelisable

Time complexity

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

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

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



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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 Results

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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 ω(v, w) are stored with the edges, for all v ∈ V_s and w ∈ V_s ∪ H_s.
- E_s = {(v, w) ∈ E : v ∈ V_s} is the subset of all the edges connected to V_s.

Outline

atching

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion



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



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

```
Clustering
Introduction
Sequential
GPU algorithm
Results
```

2D partitioning

2D matching

Conclusion

References



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 $P(\phi(pref(v)))$

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion



How to propose



Source: www.theguardian.com

proposal(v, w): v proposes to w



Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Parallel algorithm for P(s): main loop

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

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Parallel algorithm for P(s): communication

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 $Adi_{\star} := Adi_{\star} \setminus \{v\}$ SetNewPreference(x, Adj_x , pref, V_s , D_s , M_s , ϕ)

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

```
Clustering
Introduction
Sequential
GPU algorithm
Results
```

2D partitioning

2D matching

Conclusion

References

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 Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



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 algorithn Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion



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?

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm
- GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Sending messages

- The BSP system takes care that messages are sent automatically, in bulk. A useful BSPlib primitive for doing this is bsp_send.
- In the next superstep, all received messages are read (using bsp_move) and processed.
- Google's Pregel system (Malewicz 2010) follows this BSP style.

Outline

- Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results
- Clustering
- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



MulticoreBSP enables shared-memory BSP

Multicore	BSP efficient,	easy & correct parallelis	ation	ULeuven Flanders ExeScien	ce Lab Utrecht University
	HOME SOFTW	VARE COMMUNITY FORU	м		
Introduction # BSP model	Home $ ightarrow$ Introduction				Print this page
BSP library	RECENT I Version Multico see all news	NEWS n 1.2, and a roadmap for future rel reBSP for C, version 1.1 released. s items	38565 .		
	Introductio	n			
	MulticoreBSP brings B performance codes:	Speed of SpMV	multiplication or	ullicore processors. BSP program	nming leads to high-
		8 -		Cilk CSB11 PThread 1D BSP 2D	
	G£lop/s	4			
		2 0 Freescale1 ldoc	r cage15	adaptive wiki07	
Albert	-Jan Yz	elman 2014	1. www.i	nulticoreb	sp.org 🂐

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

Matching with MulticoreBSP

- BSP program can remain the same, giving portability.
- To exploit the ease of reading data in shared memory, the bsp_direct_get is available in MulticoreBSP.
- This performs the communication immediately and blocks until the communication has been carried out.
- Possible use: replace the set M_s of matched edges by a boolean array matched_s marking the local matched vertices.
- ► This array can be read by all processors using bsp_direct_get, to replace the check (x, pref(x)) ∉ M_s.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

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.

Universiteit Utrecht

Outlin

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

GPU implementation

- The graph (neighbour ranges, indices, and weights) is stored as a triplet of 1D textures (read-only arrays).
- ▶ We create one thread for each vertex in V.
- Each vertex $v \in V$ only updates
 - its colour/matching value $\pi(v)$;
 - and its proposal/response value $\sigma(v)$.
- $\pi(v) = \pi(w)$ means $(v, w) \in M$.
- Both π and σ are stored in 1D arrays in global memory. and hence are visible to all threads.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References





Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ 2

Greedv

GPU algorithm

Sequential

Results



Greedv

GPU algorithm

Sequential

Results



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

- Introduction Sequential GPU algorithm Results
- 2D partitioning
- 2D matching
- Conclusion
- References

Quality of the matching



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



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

Random colouring of the vertices

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

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

- 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.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm **GPU algorithm** Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Choosing the probability p



Following the expectation formula, we should choose $p \approx 0.53406$.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Experimental results (Fagginger Auer & Bisseling 2012)

- Implementation on the GPU using CUDA, on the CPU using Intel Threading Building Blocks (TBB).
- We consider both greedy random and greedy weighted matching.
- Test set: 10th DIMACS challenge on graph partitioning and University of Florida Sparse Matrix Collection.
- Test hardware: dual quad-core Xeon E5620 and an NVIDIA Tesla C2050 (the Little Green Machine).

Outline

Matching Introductior

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion



Results: strong scaling



Scaling of Intel TBB implementation (8 physical cores + hyperthreading).

Universiteit Utrecht

Results

Sequential
Parallel vs. sequential greedy random matching



Matching size and speedup.



Universiteit Utrecht

(日) (部) (音) (音) (音) (音) (の)

Parallel vs. sequential greedy weighted matching



Matching weight and speedup.



Parallel weighted vs. sequential greedy matching



Matching weight and speedup. Sequential greedy matching is 1/2-approximation.



Clustering of road network of the Netherlands





(f) Best clustering (G^{21})

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 Results

Clustering

Introduction

Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

DIMACS challenge February 2012



As stated on their main website, the "DIMACS Implementation Challenges address questions of determining realistic algorithm performance where worst case analysis is overly pessimistic and probabilistic models are too unrealistic: experimentation can provide guides to realistic algorithm performance where analysis fails."

For the 10th DIMACS Implementation Challenge, the two related problems of graph partitioning and graph clustering were chosen. Graph partitioning and graph clustering are among the aforementioned questions or problem areas where theoretical and practical results deviate significantly from each other, so that experimental outcomes are of particular interest.

News

- Dec 27, 2011: Tim Mattson (Intel) has joined the advisory board. Welcome, Tim!
- Dec 1 and 12, 2011: Updated workshop registration information, new graphs in category numerical!

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction

Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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.
- ► Extreme cases: a single large cluster, |V| single-vertex clusters.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction

Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



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

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

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

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction

Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Merging clusters: change in modularity

The weight of a cluster is

$$\zeta(\mathcal{C}) = \sum_{\mathbf{v}\in\mathcal{C}}\zeta(\mathbf{v}).$$

▶ 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 the modularity of the new clustering C' is

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

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

Sequential

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 GPU algori Results

Parallelisation

- Based on slight adaptations of functions from Thrust, an open-source template library for developing CUDA applications (modelled after C++ STL).
- Also, for... parallel do constructs indicating a for-loop where each iteration can be executed in parallel.

Outline

Matching

Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Results: clustering time for DIMACS graphs



- 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

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

Results: strong scaling



The clustering time as a function of the number of threads.

- Graphs from the category random/ with $2^{15}-2^{24}$ vertices.
- Intel TBB implementation on 2 quad-core 2.4 GHz Intel Xeon E5620 processors with up to 16 threads by hyperthreading.

Universiteit Utrecht

Results

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
italy	6,686,493	7,013,978	1.00	1.54	1.00	5.26 GPU algorithm
great-britain	7,733,822	8,156,517	1.00	1.79	1.00	6.00 Clustering
germany	11,548,845	12,369,181	1.00	2.82	1.00	9.57 Introduction
asia	11,950,757	12,711,603	1.00	2.69	1.00	9.33 Sequential GPU algorithm
europe	50,912,018	54,054,660	-		1.00	45.21 Results
coAuthorsCite	227,320	814,134	0.84	0.42	0.85	0.23 ^{2D partitioning}
coAuthorsDBLP	299,067	977,676	0.75	0.59	0.76	0.28 ^{2D matching}
citationCite	268,495	1,156,647	0.64	0.89	0.68	0.32 Conclusion
coPapersDBLP	540,486	15,245,729	0.64	6.43	0.67	2.28 References
coPapersCite	434,102	16,036,720	0.75	6.49	0.77	2.27

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



Sparse matrix partitioning and graph partitioning

A sparse matrix is the adjacency matrix of a sparse graph:

$$a_{ij} \neq 0 \Leftrightarrow (i,j) \in E$$

- Partitioning the nonzeros of a matrix is the same as partioning the edges of a graph.
- 2D partitioning splits both rows and columns.
- Partitioning for parallel sparse matrix-vector multiplication (SpMV) can be used in Google PageRank computation.
- Partitioning for SpMV also gives a good partitioning for many graph computations.

Outline

Aatching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching Conclusion



Advantage of 2D partitioning

- We can use both dimensions of the matrix to reduce SpMV communication.
- For a √p × √p block distribution, each matrix row or column is distributed over at most √p processors, instead of p processors for a 1D distribution.
- Relatively dense rows and columns can be split and do not cause load imbalance or memory overflow.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching Conclusion



Methods for 2D partitioning

(日) (部) (注) (注) (三)

- Existing 2D methods: coarse-grain, fine-grain, Mondriaan.
- ► New medium-grain method (Pelt & Bisseling 2014) based on splitting the m × n matrix

$$A=A^r+A^c,$$

putting a nonzero a_{ij} into A^r if row *i* has less nonzeros than column *j*, and in A^c otherwise.

► Then partition the (m + n) × (m + n) matrix B by a 1D column partitioning:

$$B = \left[\begin{array}{cc} I_n & (A^r)^T \\ A^c & I_m \end{array} \right],$$

where I_m is the identity matrix of size $m \times m$.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching Conclusion

References

Pajek Graph Drawing contest 1997





Source: University of Florida Sparse Matrix Collection



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

- 2D matching
- Conclusion
- References



Medium-grain method for partitioning



- \blacktriangleright 47 \times 47 matrix gd97_b with 264 nonzeros
- Partitioning for 2 processors
- Communication volume = 11, which is optimal

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering Introduction Sequential GPU algorithm Results

2D partitioning

- 2D matching
- Conclusion
- References



Communication volume for 2 processors



- ▶ Test set: 2264 Florida matrices, $500 \le nz \le 5,000,000$
- LB = localbest (original Mondriaan) = best of 1D row and 1D column partitioning
- ▶ FG = fine-grain
- MG = medium-grain
- IR = iterative refinement, to improve the partitioning

Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



2D (edge-based) parallel matching

	SpMV		Matching		
Name	1D	2 <i>D</i>	1D	2 <i>D</i>	
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	
$sw1~({ m small}~{ m world})$	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	

Sequential

2D matching

Communication volume in sparse matrix-vector multiplication and Karp-Sipser matching. Source: Patwary, Bisseling, Manne (2010). ・ロト ・四ト ・ヨト ・ヨト -2



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;
 - 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 Europe's road network with 51M vertices and 54M edges in 45 seconds on an 8-core CPU.
- Partitioning for sparse matrix-vector multiplication reduces communication volume for Karp–Sipser matching as well:

$$rac{1}{2}$$
 Vol(SpMV) \leq Vol(Matching) $\leq rac{3}{2}$ Vol(SpMV).

 Parallel graph algorithms will benefit from partitioning the edges instead of the vertices.



Universiteit Utrecht

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

It's all about the connections...



Outline

Matching Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

Merci beaucoup!



Further reading I

 Rob H. Bisseling, Bas O. Fagginger Auer, A. N. Yzelman, Tristan van Leeuwen, and Ümit V. Çatalyürek.
 Two-dimensional approaches to sparse matrix partitioning. In Uwe Naumann and Olaf Schenk, editors, *Combinatorial Scientific Computing*, Computational Science Series, pages 321–349. CRC Press, Taylor & Francis Group, Boca Raton, FL, 2012.

 Bas O. Fagginger Auer and Rob H. Bisseling.
 A GPU algorithm for greedy graph matching.
 In Rainer Keller, David Kramer, and Jan-Philipp Weiss, editors, Proceedings Facing the Multicore Challenge II, Karlsruhe 2011, volume 7174 of Lecture Notes in Computer Science, pages 108–119. Springer-Verlag, Berlin, 2012.

Universiteit Utrecht

Sequential

Further reading II

Bas O. Fagginger Auer and Rob H. Bisseling. Graph coarsening and clustering on the GPU. In David A. Bader, Henning Meyerhenke, Peter Sanders, and Dorothea Wagner, editors, *Graph Partitioning and Graph Clustering*, volume 588 of *Contemporary Mathematics*, pages 223–240. AMS, Providence, RI, 2013.

Fredrik Manne and Rob H. Bisseling.

A parallel approximation algorithm for the weighted maximum matching problem.

In Proceedings Seventh International Conference on Parallel Processing and Applied Mathematics (PPAM 2007), volume 4967 of Lecture Notes in Computer Science, pages 708–717, 2008.



Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Further reading III

Md. Mostofa Ali Patwary, Rob H. Bisseling, and Fredrik Manne.

Parallel greedy graph matching using an edge partitioning approach.

In Proceedings of the fourth international workshop on High-level parallel programming and applications, HLPP '10, pages 45–54, New York, NY, USA, 2010. ACM.

Daniël M. Pelt and Rob H. Bisseling. A medium-grain method for fast 2D bipartitioning of sparse matrices.

In Proceedings IEEE International Parallel and Distributed Processing Symposium 2014. IEEE Press, 2014.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References



Further reading IV

A. N. Yzelman, R. H. Bisseling, D. Roose, and
K. Meerbergen.
MulticoreBSP for C: a high-performance library for shared-memory parallel programming.
International Journal of Parallel Programming, 2013.

Outline

Matching

Introduction Greedy Parallelisable BSP algorithm GPU algorithm Results

Clustering

Introduction Sequential GPU algorithm Results

2D partitioning

2D matching

Conclusion

References

