Load balancing and random graphs

Nick Wormald
University of Waterloo

Load balancing and random graphs

Nick Wormald
University of Waterloo

Includes joint work with Jane Gao

Throw n balls randomly into n bins.

Throw n balls randomly into n bins.

Exercise: Max number in a bin is $\sim \log n / \log \log n$ (a.a.s.)

Throw n balls randomly into n bins.

Exercise: Max number in a bin is $\sim \log n / \log \log n$ (a.a.s.)

Theorem [Azar, Broder, Karlin and Upfal, '94]

Throw the balls sequentially, each ball put into the least-full of $h \geq 2$ randomly chosen bins.
Then finally, max is $O(\ln \ln n / \ln h)$.

An application

Balance the load of machines, on-line.

'load' equals number of jobs assigned.

h choices for each job.

Then the max 'load' is about $(\ln \ln n)/\ln h$.

An application

Balance the load of machines, on-line.

'load' equals number of jobs assigned.

h choices for each job.

Then the max 'load' is about $(\ln \ln n)/\ln h$.

Theorem [ABKU]

With m jobs and n machines, $m = O(n)$,
max load is a.a.s. at most $(\ln \ln n)/\ln h + O(m/n)$.

An application

Balance the load of machines, on-line.

'load' equals number of jobs assigned.

h choices for each job.

Then the max 'load' is about $(\ln \ln n)/\ln h$.

Theorem [ABKU]

With m jobs and n machines, $m = O(n)$,
max load is a.a.s. at most $(\ln \ln n)/\ln h + O(m/n)$.

Theorem [Berenbrink, Czumaj, Steger & Vöcking, '00]

With arbitrary m,
max load is a.a.s. $m/n + O(\ln \ln n)$ (fixed $h \geq 2$).

Another application: Disk allocation

Large array of n disks.

Another application: Disk allocation

Large array of n disks.

m read requests arrive.

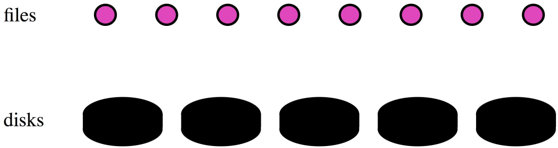

Another application: Disk allocation

Large array of n disks.

m read requests arrive.

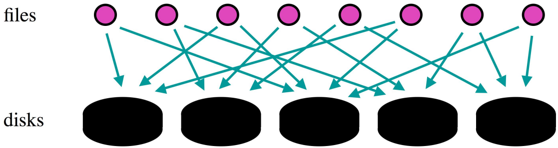

Each file can choose between two disks.

Another application: Disk allocation

Large array of n disks.

m read requests arrive.

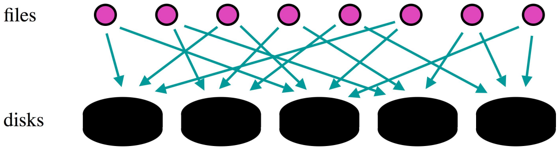

Each file can choose between two disks.

Task: Balance the loads.

Backlog of processing $\rightarrow$ off-line load balancing

Backlog of processing → off-line load balancing

Efficient ($O(m^2)$) optimal algorithm exists — solving a max flow problem.

Questions

(i) (Threshold) Fix k. What is the threshold m at which the maximum load will first exceed k.

Questions

(i) (Threshold) Fix k. What is the threshold m at which the maximum load will first exceed k.

(ii) (Approximate algorithms) What about faster algorithms that are ‘nearly optimal’, i.e. give the optimal load for most inputs with given parameters?

Offline load balancing: a graph orientation problem

Offline load balancing: a graph orientation problem

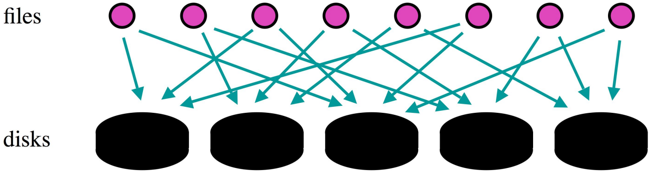

Offline load balancing: a graph orientation problem

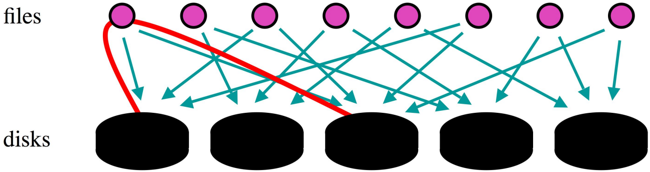

Offline load balancing: a graph orientation problem

files =
edges

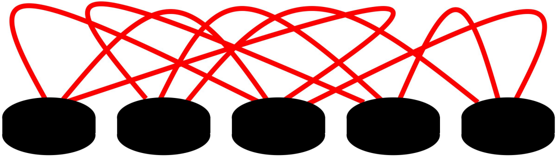

disks =
vertices

Offline load balancing: a graph orientation problem

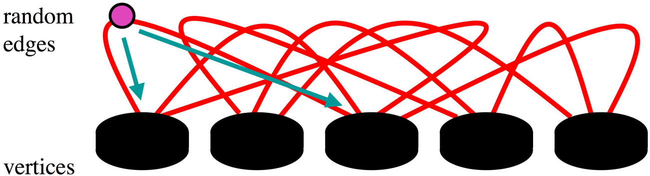

Offline load balancing: a graph orientation problem

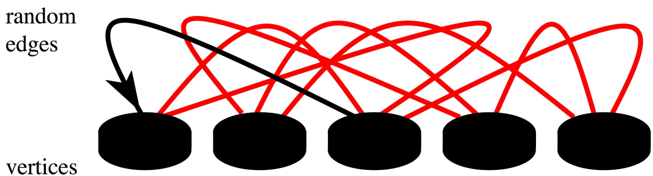

Offline load balancing: a graph orientation problem

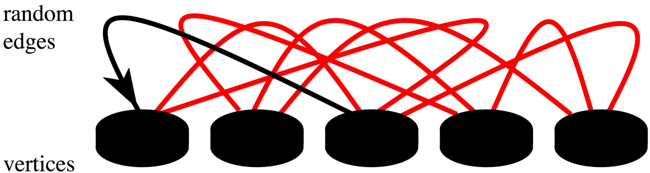

A graph G is *k*-orientable if the edges can be oriented so that every vertex has indegree at most k.

Offline load balancing: a graph orientation problem

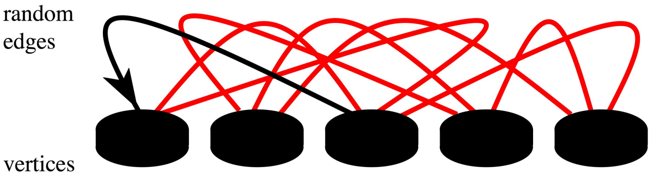

A graph G is *k*-orientable if the edges can be oriented so that every vertex has indegree at most k.

This condition holds iff max load is k.

Offline load balancing: a graph orientation problem

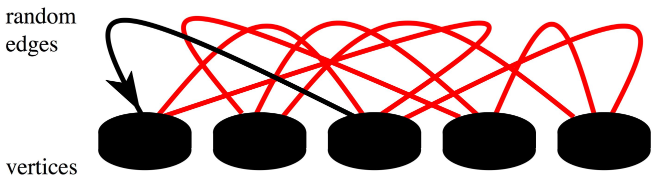

A graph G is *k*-orientable if the edges can be oriented so that every vertex has indegree at most k.

This condition holds iff max load is k.

Assumption: we consider G random in the space $\mathcal{G}(n, m)$.

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

$k = 2$:

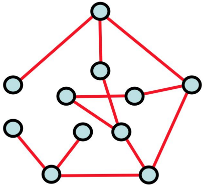

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

$k = 2$:

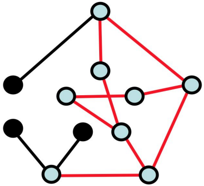

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

$k = 2$:

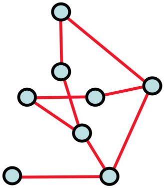

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

$k = 2$:

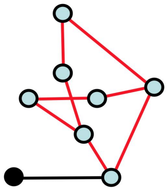

Cores and orientability

The k-core of a graph is the largest subgraph with all vertices of degree at least k.

$k = 2$:

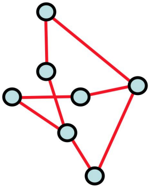

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k + 1)$-core, it is k-orientable.

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k+1)$-core, it is k-orientable.

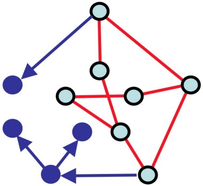

Cores and orientability

The *k-core* of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k + 1)$-core, it is k-orientable.

The converse is not true.

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k+1)$-core, it is k-orientable.

The converse is not true.

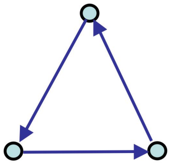

Cores and orientability

The k-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k+1)$-core, it is k-orientable.

The converse is not true.

If G has a subgraph with average degree greater than $2k$ then G is not k-orientable.

Cores and orientability

The *k*-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k+1)$-core, it is k-orientable.

The converse is not true.

If G has a subgraph with average degree greater than $2k$ then G is not k-orientable.

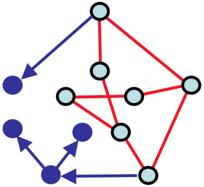

Cores and orientability

The k-core of a graph is the largest subgraph with all vertices of degree at least k.

Clearly, if a graph has no $(k+1)$-core, it is k-orientable.

The converse is not true.

If G has a subgraph with average degree greater than $2k$ then G is not k-orientable.

The converse is also true(!)

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

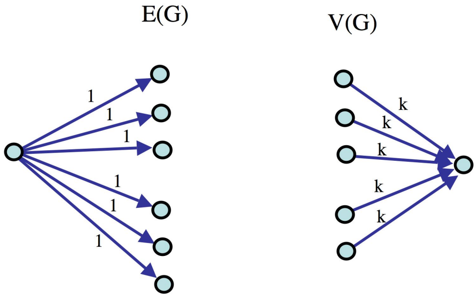

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

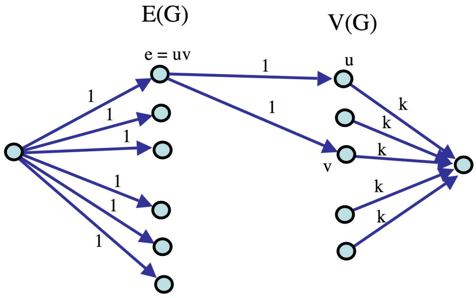

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

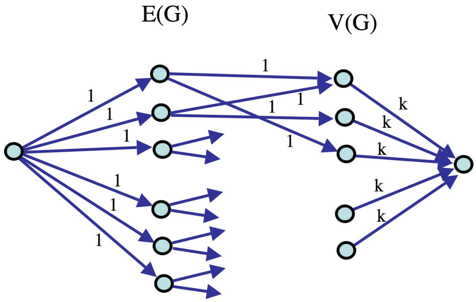

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

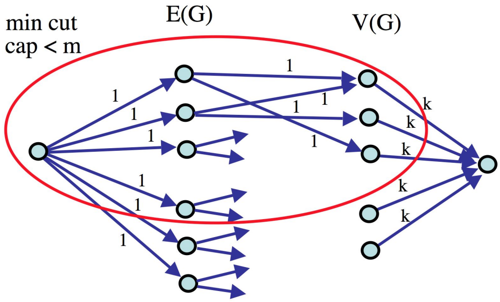

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

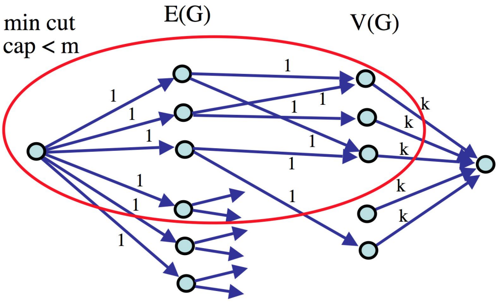

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

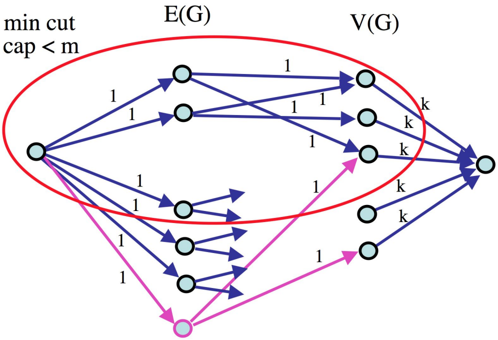

Theorem [Hakimi, '65]

G is not k-orientable implies G has a subgraph with average degree greater than $2k$.

Proof If not k-orientable then the following has max flow of value less than $m = |E(G)|$.

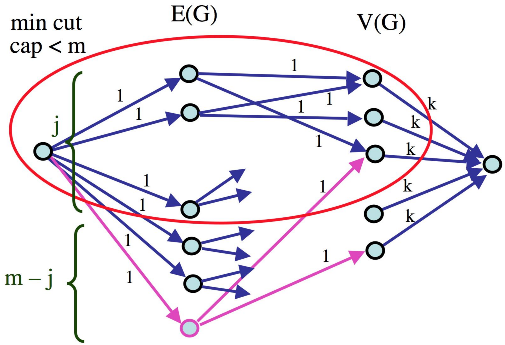

Theorem [Hakimi, '65]

G is not *k*-orientable implies *G* has a subgraph with average degree greater than $2k$.

Proof If not *k*-orientable then the following has max flow of value less than $m = |E(G)|$.

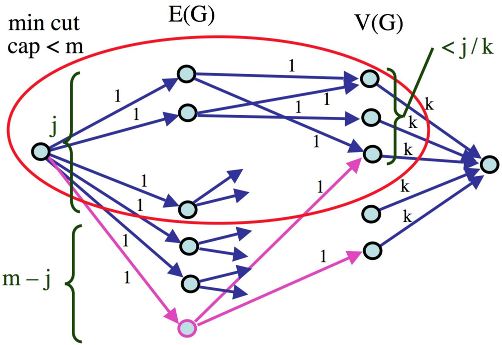

Old question

What is the threshold for $G \in \mathcal{G}(n, m)$ possessing a k-core?

Old question

What is the threshold for $G \in \mathcal{G}(n, m)$ possessing a k-core?

Theorem [Pittel, Spencer, W. '95]

The threshold for having a k-core is $m = c_k n/2$ where c_k is a constant depending on k.

Old question

What is the threshold for $G \in \mathcal{G}(n, m)$ possessing a k-core?

Theorem [Pittel, Spencer, W. '95]

The threshold for having a k-core is $m = c_k n/2$ where c_k is a constant depending on k.

$c_k = k + d_k\sqrt{k} + O(1)$.

Newer question

What is the threshold for k-orientability of $G \in \mathcal{G}(n, m)$?

Newer question

What is the threshold for k-orientability of $G \in \mathcal{G}(n, m)$?

Karp and Saks ['98] conjectured that it is the same as the threshold for

$$\bar{d}_{k+1} > 2k$$

where $\bar{d}_{k+1}(G)$ is the average degree of vertices in the $(k+1)$-core of G.

Newer question

What is the threshold for k-orientability of $G \in \mathcal{G}(n, m)$?

Karp and Saks ['98] conjectured that it is the same as the threshold for

$$\bar{d}_{k+1} > 2k$$

where $\bar{d}_{k+1}(G)$ is the average degree of vertices in the $(k+1)$-core of G.

This was proved in '07 by Cain, Sanders and Wormald, and independently by Fernholz and Ramachandran.

Newer question

What is the threshold for k-orientability of $G \in \mathcal{G}(n, m)$?

Karp and Saks ['98] conjectured that it is the same as the threshold for

$$\bar{d}_{k+1} > 2k$$

where $\bar{d}_{k+1}(G)$ is the average degree of vertices in the $(k+1)$-core of G.

This was proved in '07 by Cain, Sanders and Wormald, and independently by Fernholz and Ramachandran.

This threshold for m is $\mu_{k+1} n/2$ where μ_{k+1} is constant.

$m < (\mu_{k+1} - \epsilon)n/2$ implies k-orientable (a.a.s.)

Proved algorithmically in both cases doing orientations edge by edge.

Note: we always orient the remaining incident edges to a vertex if doing so does not exceed load k.

$m < (\mu_{k+1} - \epsilon)n/2$ implies k-orientable (a.a.s.)

Proved algorithmically in both cases doing orientations edge by edge.

Note: we always orient the remaining incident edges to a vertex if doing so does not exceed load k.

load-degree algorithm

* orient edges one by one to vertices.
* a vertex with j edges oriented to it and i unoriented has 'load degree' $j + i/2$. Greedily choose a vertex of minimum load degree and orient an unoriented edge to it.

$m < (\mu_{k+1} - \epsilon)n/2$ implies k-orientable (a.a.s.)

Proved algorithmically in both cases doing orientations edge by edge.

Note: we always orient the remaining incident edges to a vertex if doing so does not exceed load k.

load-degree algorithm

* orient edges one by one to vertices.
* a vertex with j edges oriented to it and i unoriented has 'load degree' $j + i/2$. Greedily choose a vertex of minimum load degree and orient an unoriented edge to it.

Theorem [CSW '07]

For $c < \mu_{k+1} - \epsilon$, the load-degree algorithm a.a.s. completes a k-orientation in a random graph $\mathcal{G}(n, cn/2)$.

FR's algorithm: choose a vertex of min degree, $k + j$, and do the following:

FR's algorithm: choose a vertex of min degree, $k + j$, and do the following:

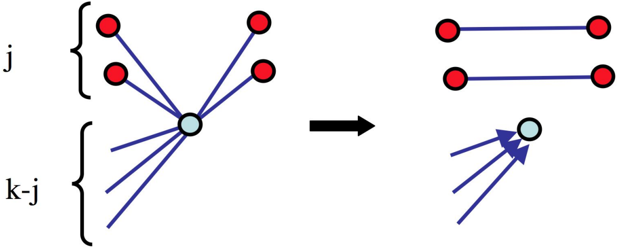

FR's algorithm: choose a vertex of min degree, $k + j$, and do the following:

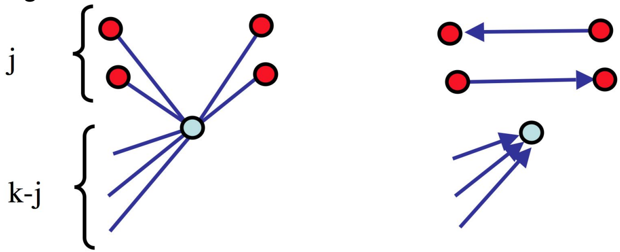

FR's algorithm: choose a vertex of min degree, $k + j$, and do the following:

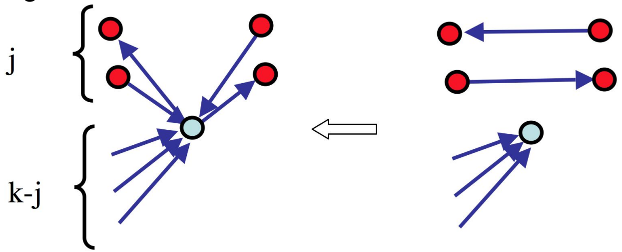

FR's algorithm: choose a vertex of min degree, $k+j$, and do the following:

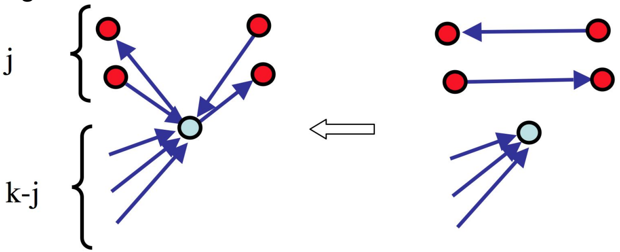

Theorem [FR]

For $c < \mu_{k+1} - \epsilon$, this algorithm a.a.s. completes a k-orientation in a random graph $\mathcal{G}(n, cn/2)$.

Problem version 2. More choices

Choosing one out of h

What if h copies of each file are stored?

Problem version 2. More choices

Choosing one out of h

What if h copies of each file are stored?

Each file can choose 1 out of h given disks.

Orientation problem becomes the following:

Given an h-uniform hypergraph G, assign ('orient') each edge of G to exactly one of its incident vertices, so that each vertex is allocated at most k edges.

Problem version 2. More choices

Choosing one out of h

What if h copies of each file are stored?

Each file can choose 1 out of h given disks.

Orientation problem becomes the following:

Given an h-uniform hypergraph G, assign ('orient') each edge of G to exactly one of its incident vertices, so that each vertex is allocated at most k edges.

The same flow proof gives:

Orientation exists iff no subhypergraph has average degree at least hk.

Clear first step: find the $(k+1)$-core.

Two approaches: extend load degree algorithm, or extend FR algorithm.

Clear first step: find the $(k+1)$-core.

Two approaches: extend load degree algorithm, or extend FR algorithm.

BOTH seem DIFFICULT!

FR:

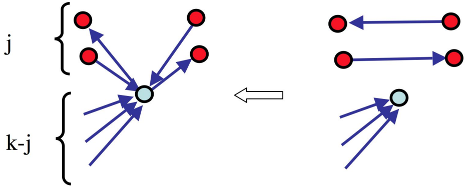

Conjecture

A.a.s. if $\bar{d} < hk - \epsilon$, where $\bar{d}$ is the average degree in the $(k+1)$-core of $G \in \mathcal{G}(n, m, h)$, then G is k-orientable.

Conjecture

A.a.s. if $\bar{d} < hk - \epsilon$, where $\bar{d}$ is the average degree in the $(k+1)$-core of $G \in \mathcal{G}(n, m, h)$, then G is k-orientable.

Theorem [Gao and W, '09]

The conjecture is true for fixed k provided it is sufficiently large.

The threshold at which this occurs (core of h-uniform hypergraph gets average degree hk) is known. (Cain and Wormald; Molloy, Cooper, Janson and Luczak.)

Problem version 3. Even more to choose from

Storing two copies of every file requires double the space.

More efficient storage strategy:

Problem version 3. Even more to choose from

Storing two copies of every file requires double the space.

More efficient storage strategy:

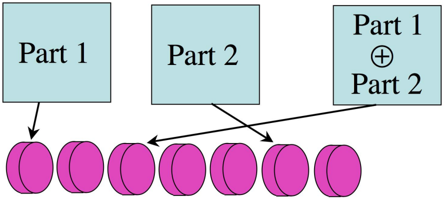

Problem version 3. Even more to choose from

Storing two copies of every file requires double the space.

More efficient storage strategy:

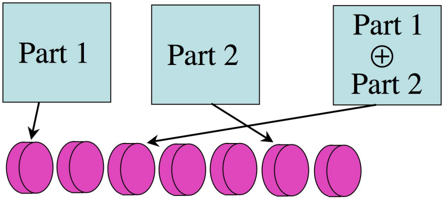

... and what if two disks fail? More general error correcting strategies.

Each file will need to choose w out of h disks, for some $1 \leq w < h$.

More general orientations

A hyperedge is w-oriented if exactly w distinct vertices in it are marked with positive signs with respect to the hyperedge.

The indegree of a vertex is the number of positive signs it receives.

A (w, k)-orientation of an h-hypergraph is a w-orientation all hyperedges such that each vertex has indegree at most k.

If such a (w, k)-orientation exists, we say the hypergraph is (w, k)-orientable .

More general orientations

A hyperedge is w-oriented if exactly w distinct vertices in it are marked with positive signs with respect to the hyperedge.

The indegree of a vertex is the number of positive signs it receives.

A (w, k)-orientation of an h-hypergraph is a w-orientation all hyperedges such that each vertex has indegree at most k.

If such a (w, k)-orientation exists, we say the hypergraph is (w, k)-orientable .

The disks can be allocated in the w-out-of-h setting iff the corresponding hypergraph is (w, k)-orientable.

More general cores

light vertex: one with degree at most k.

free step (I) : For any light vertex v, give v positive sign with respect to all incident hyperedges, delete v and remove v from all incident hyperedges *but don't delete those hyperedges* .

free step (II) : remove any hyperedge if its cardinality falls to $h - w$. (Must then have w vertices with positive signs.)

More general cores

light vertex: one with degree at most k.

free step (I) : For any light vertex v, give v positive sign with respect to all incident hyperedges, delete v and remove v from all incident hyperedges *but don't delete those hyperedges* .

free step (II) : remove any hyperedge if its cardinality falls to $h - w$. (Must then have w vertices with positive signs.)

(w, k)-core : the result of repeating free steps until no more can be taken. Has mixed edge sizes.

Conjecture

The event

$$\sum_{j=0}^{w-1}(w-j)m_{h-j} < k\bar{n},$$

where

m_{h-j} *is the number of edges of cardinality* $h-j$ *in the* $(w, k+1)$*-core of* $G \in \mathcal{G}(n, m, h)$*,*

$\bar{n}$ *is its number of vertices,*

has the same sharp threshold as G *being* (w, k)*-orientable.*

Conjecture

The event

$$\sum_{j=0}^{w-1}(w-j)m_{h-j} < k\bar{n},$$

where

m_{h-j} *is the number of edges of cardinality* $h-j$ *in the* $(w, k+1)$*-core of* $G \in \mathcal{G}(n, m, h)$*,*

$\bar{n}$ *is its number of vertices,*

has the same sharp threshold as G *being* (w, k)*-orientable. In particular* (w, k)*-orientability has a sharp threshold.*

Theorem [Gao and W, '09]

The conjecture is true for fixed k sufficiently large.

We determine the threshold in terms of the solution of a system of differential equations.

Proof approach

Look again at the flow formulation. E.g.: $h = 3$, $w = 2$.

Proof approach

Look again at the flow formulation. E.g.: $h = 3$, $w = 2$.

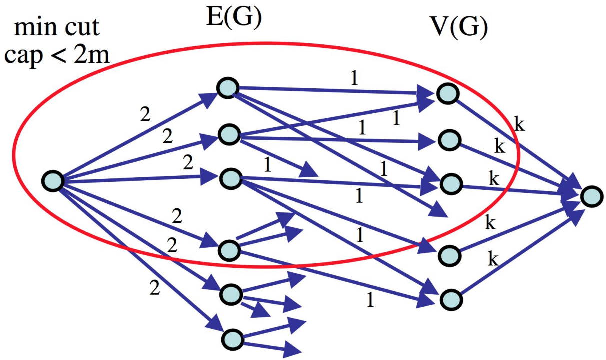

Proof approach

Look again at the flow formulation. E.g.: $h = 3$, $w = 2$.

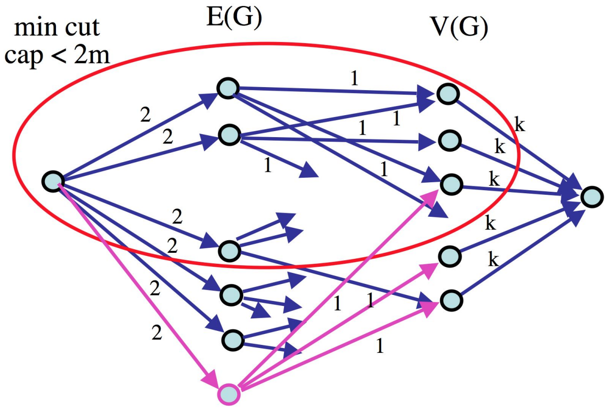

From the flow:

Lemma

An h-hypergraph G can be (w, k)-oriented iff there is no set $S \subseteq V(G)$ with

$$\frac{1}{|S|} \sum_{x \in E(G)} \left(|x \cap S| - (h - w)\right)^{+} > k.$$

From the flow:

Lemma

An h-hypergraph G can be (w, k)-oriented iff there is no set $S \subseteq V(G)$ with

$$\frac{1}{|S|} \sum_{x \in E(G)} \left(|x \cap S| - (h - w)\right)^{+} > k.$$

Major result:

Sub-theorem

Let $\mathcal{C}$ denote the $(w, k+1)$-core of $G \in \mathcal{G}(n, m, h)$. A.a.s. if

$$\frac{1}{|S|} \sum_{x \in E(G)} \left(|x \cap S| - (h - w)\right)^{+} < k - \epsilon$$

for $S = \mathcal{C}$ then it holds for all $S \subseteq V(G)$.

For $S = \mathcal{C}$, the requirement becomes

$$\frac{1}{|\mathcal{C}|} \sum_{j=0}^{w-1} (w - j) m_{h-j} < k - \epsilon$$

where m_{h-j} is the number of edges of the core of size $h - j$. This gives the theorem.

If $w = 1$ this says the average degree is less than $h(k - \epsilon)$.

For $S = \mathcal{C}$, the requirement becomes

$$\frac{1}{|\mathcal{C}|}\sum_{j=0}^{w-1}(w-j)m_{h-j} < k - \epsilon$$

where m_{h-j} is the number of edges of the core of size $h-j$. This gives the theorem.

If $w = 1$ this says the average degree is less than $h(k-\epsilon)$.

To prove the sub-theorem we need properties of the random $(w, k+1)$-core.

For $S = \mathcal{C}$, the requirement becomes

$$\frac{1}{|\mathcal{C}|}\sum_{j=0}^{w-1}(w-j)m_{h-j} < k - \epsilon$$

where m_{h-j} is the number of edges of the core of size $h - j$. This gives the theorem.

If $w = 1$ this says the average degree is less than $h(k - \epsilon)$.

To prove the sub-theorem we need properties of the random $(w, k+1)$-core.

To study the threshold we also need properties of the random $(w, k+1)$-core.

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \leq i \leq m)$.

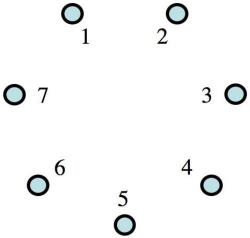

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \leq i \leq m)$.

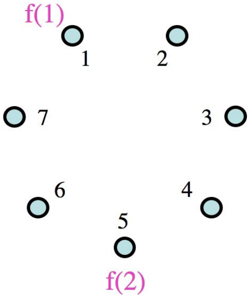

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \leq i \leq m)$.

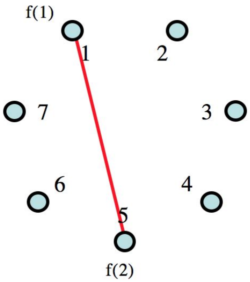

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \leq i \leq m)$.

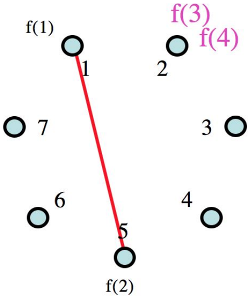

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \le i \le m)$.

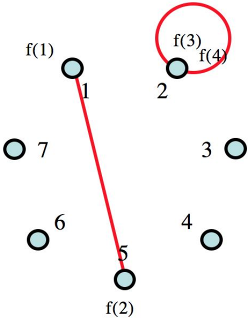

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \leq i \leq m)$.

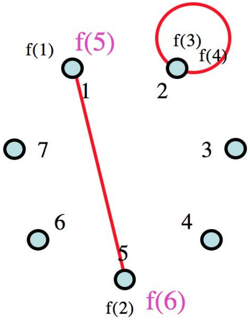

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \le i \le m)$.

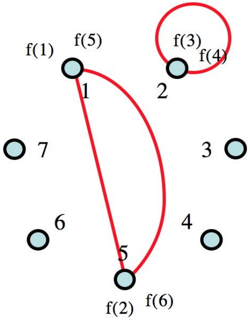

Analysing the random $(w, k+1)$-core

A random pseudograph model of Bollobás and Frieze ('85), Chvátal ('91), etc:

Given positive integers n and m, define the probability space of functions $f : [2m] \to [n]$, all functions equiprobable.

Pseudograph $\mathcal{P}(n, m)$: edges are $\{f(2i-1), f(2i)\}$ $(1 \le i \le m)$.

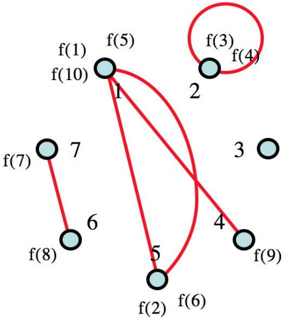

Important properties of the pseudograph model

Suppose that $m = cn/2$.

- Simple graphs occur uniformly at random. That is, $\mathcal{P}(n, m)$ restricted to simple graphs is $\mathcal{G}(n, m)$.

Important properties of the pseudograph model

Suppose that $m = cn/2$.

- Simple graphs occur uniformly at random. That is, $\mathcal{P}(n, m)$ restricted to simple graphs is $\mathcal{G}(n, m)$.
- $\mathbf{P}(\mathcal{P}(n, m) \text{ is simple}) \sim e^{-c/2-c^2/4}$.

Important properties of the pseudograph model

Suppose that $m = cn/2$.

- Simple graphs occur uniformly at random. That is, $\mathcal{P}(n, m)$ restricted to simple graphs is $\mathcal{G}(n, m)$.
- $\mathbf{P}(\mathcal{P}(n, m) \text{ is simple}) \sim e^{-c/2-c^2/4}$.
- Connection property: events a.a.s. true for $\mathcal{P}(n, m)$ (c fixed) are a.a.s. true for $\mathcal{G}(n, m)$.

Important properties of the pseudograph model

Suppose that $m = cn/2$.

- Simple graphs occur uniformly at random. That is, $\mathcal{P}(n, m)$ restricted to simple graphs is $\mathcal{G}(n, m)$.
- $\mathbf{P}(\mathcal{P}(n, m) \text{ is simple}) \sim e^{-c/2-c^2/4}$.
- Connection property: events a.a.s. true for $\mathcal{P}(n, m)$ (c fixed) are a.a.s. true for $\mathcal{G}(n, m)$.
- Multinomial property of degree sequence $(X_1, \ldots, X_n)$: for nonnegative integer vector $(d_1, \ldots, d_n)$ with $\sum_i d_i = 2m$,

$$\mathbf{P}(X_i = d_i \text{ for } 1 \leq i \leq n) = \frac{(2m)!}{n^{2m} \prod_{i=1}^{n} d_i!}.$$

Our adaptation

For a random pseudo-hypergraph $\mathcal{P}_n(h_1, \ldots, h_m)$ with

- n vertices
- m edges of cardinalities $h_1, \ldots, h_m$

Change domain of the random mapping from $1, \ldots, 2m$ to a set of balls with h_i of them labelled i ($1 \leq i \leq m$).

Our adaptation

For a random pseudo-hypergraph $\mathcal{P}_n(h_1, \ldots, h_m)$ with

- n vertices
- m edges of cardinalities $h_1, \ldots, h_m$

Change domain of the random mapping from $1, \ldots, 2m$ to a set of balls with h_i of them labelled i $(1 \leq i \leq m)$.

Finding the $(w, k+1)$-core of the pseudo-hypergraph

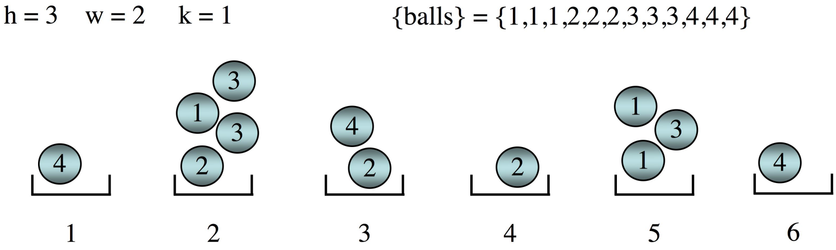

Finding the $(w, k+1)$-core of the pseudo-hypergraph

h = 3 w = 2 k + 1 = 2

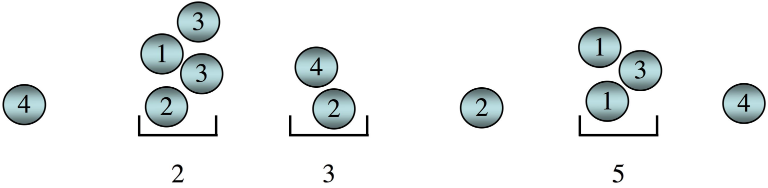

Finding the $(w, k+1)$-core of the pseudo-hypergraph

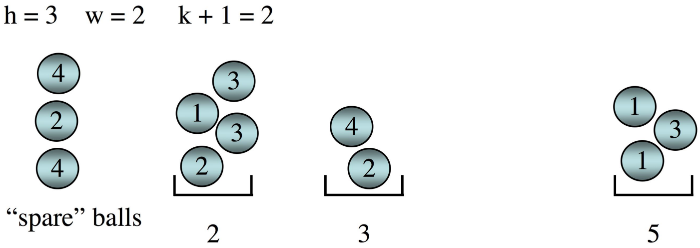

Finding the $(w, k+1)$-core of the pseudo-hypergraph

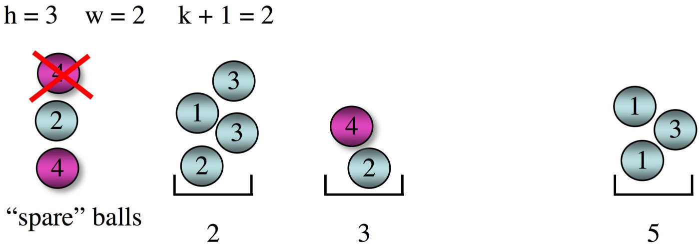

Finding the $(w, k+1)$-core of the pseudo-hypergraph

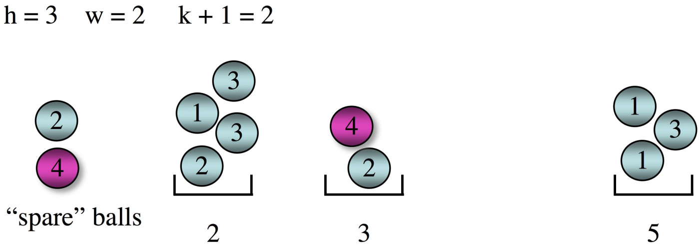

Finding the $(w, k+1)$-core of the pseudo-hypergraph

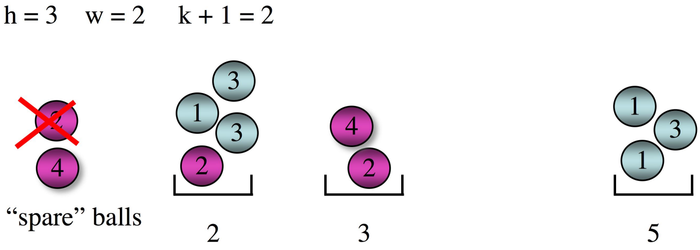

Finding the $(w, k+1)$-core of the pseudo-hypergraph

h = 3 w = 2 k + 1 = 2

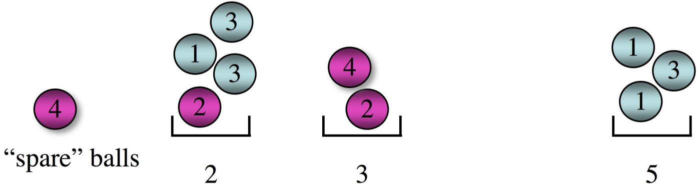

Finding the $(w, k+1)$-core of the pseudo-hypergraph

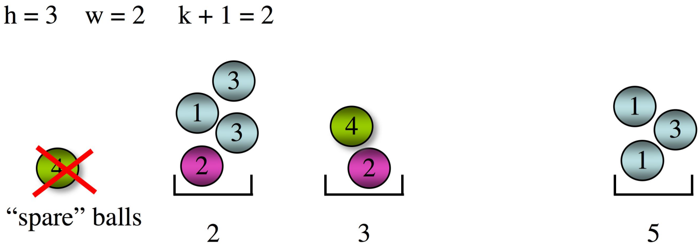

Finding the $(w, k+1)$-core of the pseudo-hypergraph

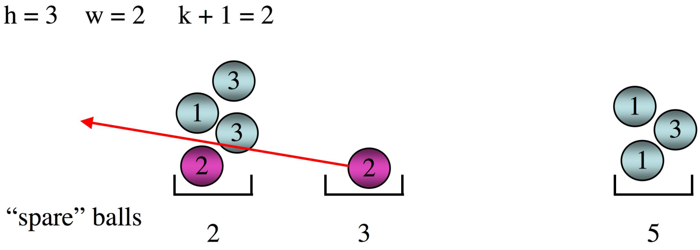

Finding the $(w, k+1)$-core of the pseudo-hypergraph

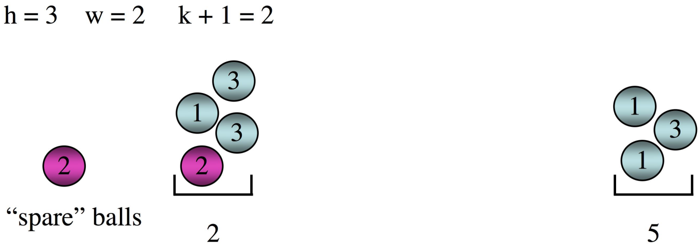

Finding the $(w, k+1)$-core of the pseudo-hypergraph

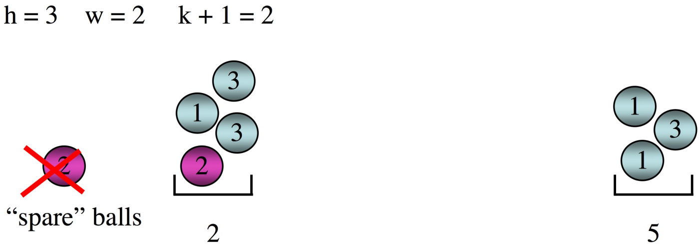

Finding the $(w, k+1)$-core of the pseudo-hypergraph

$h = 3 \quad w = 2 \quad k + 1 = 2$

Analysis

At each step, the remaining bins are filled with a random allocation, conditional upon the number of balls of each label and that each bin receives at least $k + 1$ balls.

Analysis

At each step, the remaining bins are filled with a random allocation, conditional upon the number of balls of each label and that each bin receives at least $k + 1$ balls.

The number of balls in the bins therefore has a 'truncated' multinomial distribution which is close to a set of copies of truncated Poisson:

$$\mathbf{P}(X = j) = \frac{e^{-\lambda}\lambda^j/j!}{\sum_{i \geq k+1} e^{-\lambda}\lambda^i/i!}.$$

Asymptotic Poisson distribution helps determine the expected changes in numbers of edges of each cardinality at each step.

Asymptotic Poisson distribution helps determine the expected changes in numbers of edges of each cardinality at each step.

The differential equation method shows sharp concentration for the final core (solve numerically).

Asymptotic Poisson distribution helps determine the expected changes in numbers of edges of each cardinality at each step.

The differential equation method shows sharp concentration for the final core (solve numerically).

Even without d.e.'s we can prove (after considerable work, and for large k)

Sub-theorem

Let $\mathcal{C}$ denote the $(w, k+1)$-core of $G \in \mathcal{G}(n, m, h)$. A.a.s. if

$$\frac{1}{|S|} \sum_{x \in E(G)} \left(|x \cap S| - (h - w)\right)^{+} < k - \epsilon$$

for $S = \mathcal{C}$ then it holds for all $S \subseteq V(G)$.

Summary

- For large k we have shown that k-orienting the random h-uniform hypergraph has a threshold determined by the density of the $(k+1)$-core. This generalises (for large k) the result known for graphs that proved Karp and Saks' conjecture, with a new simpler proof.

Summary

- For large k we have shown that k-orienting the random h-uniform hypergraph has a threshold determined by the density of the $(k+1)$-core. This generalises (for large k) the result known for graphs that proved Karp and Saks' conjecture, with a new simpler proof.
- Even further generalising, existence of a (w, k)-orientation (k large) of the random h-uniform hypergraph is essentially determined by the existence of a subgraph with a certain 'density'.

Summary

- For large k we have shown that k-orienting the random h-uniform hypergraph has a threshold determined by the density of the $(k+1)$-core. This generalises (for large k) the result known for graphs that proved Karp and Saks' conjecture, with a new simpler proof.
- Even further generalising, existence of a (w, k)-orientation (k large) of the random h-uniform hypergraph is essentially determined by the existence of a subgraph with a certain 'density'.
- Conjecture the same results for all $k \geq 1$.

Summary

- For large k we have shown that k-orienting the random h-uniform hypergraph has a threshold determined by the density of the $(k+1)$-core. This generalises (for large k) the result known for graphs that proved Karp and Saks' conjecture, with a new simpler proof.
- Even further generalising, existence of a (w, k)-orientation (k large) of the random h-uniform hypergraph is essentially determined by the existence of a subgraph with a certain 'density'.
- Conjecture the same results for all $k \geq 1$.
- No fast algorithm for loadbalancing yet. Conjecture that the obvious generalisation of the load-degree algorithm does the job for all (w, k).