



---

Basic Research in Computer Science

BRICS RS-98-24 D. P. Dubhashi: Talagrand's Inequality and Locality in Distributed Computing

## Talagrand's Inequality and Locality in Distributed Computing

Devdatt P. Dubhashi

BRICS Report Series

ISSN 0909-0878

RS-98-24

October 1998

**Copyright © 1998, BRICS, Department of Computer Science  
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work  
is permitted for educational or research use  
on condition that this copyright notice is  
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.  
Copies may be obtained by contacting:**

**BRICS  
Department of Computer Science  
University of Aarhus  
Ny Munkegade, building 540  
DK-8000 Aarhus C  
Denmark  
Telephone: +45 8942 3360  
Telefax: +45 8942 3255  
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide  
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`  
`ftp://ftp.brics.dk`  
**This document in subdirectory RS/98/24/**

# Talagrand's Inequality and Locality in Distributed Computing \*

Devdatt P. Dubhashi <sup>†</sup>

Department of Computer Science and Engg.

Indian Institute of Technology, Delhi

Hauz Khas, New Delhi 110016

INDIA

`emaildubhashi@cse.iitd.ernet.in`

October 6, 1998

## Abstract

We illustrate the use of Talagrand's inequality and an extension of it to dependent random variables due to Marton for the analysis of distributed randomised algorithms, specifically, for edge colouring graphs.

## 1 Introduction

The aim of this paper is to advocate the use of Talagrand's isoperimetric inequality [10] and an extension of it due to Marton [5, 6] as a tool for the analysis of distributed randomized algorithms that work in the *locality*

---

\*To appear in the Second International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM 98) Barcelona, Spain, October 8–10 1998.

<sup>†</sup>Work partly done while at BRICS, Department of Computer Science, University of Aarhus, Denmark. Partially supported by the ESPRIT Long Term Research program of the EU under contract No. 20244 (ALCOM-IT)

paradigm. Two features of the inequality are crucially used in the analysis: first, very refined control on the influence of the underlying variables can be exercised to get significantly stronger bounds by exploiting the non-uniform and asymmetric conditions required by the inequality (in contrast to previous methods) and second, the method, using an extension of the basic inequality to dependent variables due to Marton [6] succeeds in spite of lack of full independence amongst the underlying variables. This last feature especially makes it a particularly valuable tool in Computer Science contexts where lack of independence is omnipresent. Our contribution is to highlight the special relevance of the method for Computer Science applications by demonstrating its use in the context of a class of distributed computations in the locality paradigm.

We give a high probability analysis of a distributed algorithm for edge-colouring a graph [8]. Apart from its intrinsic interest as a classical combinatorial problem, and as a paradigm example for locality in distributed computing, edge colouring is also useful from a practical standpoint because of its connection to scheduling. In distributed networks or architectures an edge colouring corresponds to a set of data transfers that can be executed in parallel. So, a partition of the edges into a small number of colour classes – i.e. a “good” edge colouring – gives an efficient schedule to perform data transfers (for more details, see [8, 2]). The analysis of edge colouring algorithms published in the literature is extremely long and difficult and that in [8] is moreover, based on a certain *ad hoc* extension of the Chernoff-Hoeffding bounds. In contrast, our analysis is a very simple, short and streamlined application of Talagrand’s inequality, only two pages long.

In § 5, we outline how other edge and vertex colouring algorithms can also be tackled by the same methods in a general framework. These examples are intended moreover, as a dramatic illustration of the versatility and power of the method for the analysis of locality in distributed computing in general.

## 2 Distributed Edge Colouring

Vizing’s Theorem shows that every graph  $G$  can be edge coloured sequentially in polynomial time with  $\Delta$  or  $\Delta + 1$  colours, where  $\Delta$  is the maximum degree of the input graph (see, for instance, [1]).

It is a challenging open problem whether colourings as good as these

can be computed fast in a distributed model. In the absence of such a result one might aim at the more modest goal of computing reasonably good colourings, instead of optimal ones. By a trivial modification of a well-known *vertex* colouring algorithm of Luby it is possible to edge colour a graph using  $2\Delta - 2$  colours in  $O(\log n)$  rounds (where  $n$  is the number of processors) [4].

We shall present and analyze a simple localised distributed algorithm that compute near optimal edge colourings. The algorithm proceeds in a sequence of rounds. In each round, a simple randomised heuristic is invoked to colour a significant fraction of the edges successfully. The remaining edges are passed over to succeeding rounds. This continues until the number of edges is small enough to employ a brute-force method at the final step. For example, the algorithm of Luby mentioned above can be invoked when the degree of the graph becomes small i.e. when the condition  $\Delta \gg \log n$  is no longer satisfied.

First the algorithm invokes a reduction to bipartite graphs by a standard procedure, see [8]. We describe the action carried out by the algorithm in a single round starting with a bipartite graph. At the beginning of each round, there is a palette of fresh new available colours,  $[\Delta]$ , where  $\Delta$  is the maximum degree of the graph at the current stage.

**Algorithm P**<sup>1</sup>: There is a two step protocol:

- Each bottom vertex, in parallel, makes a *proposal* independently of other bottom vertices by assigning a random *permutation* of the colours to their incident edges.
- Each top vertex, in parallel, then picks a *winner* out of every set of incident edges that have the same colour. Tentative colours of winner edges become final.
- The *losers*- edges who are not winners- are decoloured and passed to the next round.

It is apparent that the algorithm is truly distributed. That is to say, each vertex need only exchange information with the neighbours to execute the algorithm. This and its simplicity makes the algorithms amenable for implementations in a distributed environment. Algorithm P is exactly the algorithm used in [8].

---

<sup>1</sup>for *permutation*.

We focus all our attention in the analysis of one round of the algorithm. Let  $\Delta$  denote the maximum degree of the graph at the beginning of the round and  $\Delta'$  denote the maximum degree of the leftover graph. One can easily show that  $\mathbb{E}[\Delta' \mid \Delta] \leq \beta\Delta$ , for some constant  $0 < \beta < 1$ . The goal is to show that this holds with high probability. This is done in § 4 after the relevant tools – the concentration of measure inequalities – are introduced in the next section.

### 3 Concentration of Measure Inequalities

Talagrand’s isoperimetric inequality for the concentration of measure involves a product space  $\Omega = \prod_{i \in [n]} \Omega_i$ , equipped with the product measure  $P = \prod_{i \in [n]} P_i$  (where the  $P_i$ s are arbitrary measures on the individual spaces) and, crucially, a certain notion of “convex distance” between a point  $x \in \omega$  and a subset  $A \subseteq \Omega$ :

$$d_T(x, A) := \sup_{\sum \alpha_i^2 = 1} \min_{y \in A} \sum_{x_i \neq y_i} \alpha_i. \quad (1)$$

(The sup is over all reals  $\alpha_1, \dots, \alpha_n$  satisfying the stated condition.) For sets  $A, B \subseteq \Omega$ , set  $d_T(A, B) := \min_{x \in A} d_T(x, B)$ .

Talagrand’s inequality [10] is:

$$P(A)P(B) \leq \exp(-d_T^2(A, B)/4), \quad A, B \subseteq \Omega. \quad (2)$$

In an alternative elegant approach, Marton [5, 6] shows that this inequality is in turn a consequence of certain information–theoretic inequalities. This requires an analogue to (1) for distributions. Note that for a one point set  $A = \{y\}$ , the distance (1) is

$$d_T(x, y) = \sup_{\sum \alpha_i^2 = 1} \sum_{x_i \neq y_i} \alpha_i.$$

Now, for distributions  $P$  and  $Q$  on  $\Omega$ , we define  $d_T(P, Q)$  as the minimum of  $\mathbb{E}[d_T(X, Y)]$  over all joint distributions of  $X$  and  $Y$  with marginals  $P$  and  $Q$  respectively.

We also require the notion of *conditional* or *relative entropy*  $\mathbb{H}(P \mid Q)$  of distribution  $P$  with respect to  $Q$ :

$$\mathbb{H}(P \mid Q) := \sum_{\omega \in \Omega} P(\omega) \log \frac{P(\omega)}{Q(\omega)}. \quad (3)$$

(This is also sometimes called the *informational divergence* and denoted  $D(P||Q)$ .)

The information theoretic inequality is: if  $P$  is a product measure, then for any other measure  $Q$ ,

$$d_T(P, Q) \leq \sqrt{2H(P | Q)}. \quad (4)$$

One can show that the concentration of measure inequality (2) is a consequence of (4).

Marton generalises (4) to the case when the underlying measure  $P$  is not the product measure (i.e. the corresponding variables are not independent). The setting is very similar to that in martingale inequalities where the underlying variables are “exposed” one by one. Suppose  $k \leq n$  and  $x_1^k, \hat{x}_1^k$  differ only in the last co-ordinate:

$$x_1^{k-1} = \hat{x}_1^{k-1}, x_k \neq \hat{x}_k.$$

We want to consider the corresponding conditional distributions on the remaining variables, once the first  $k$  variables are fixed to the values  $x_1^k$  and  $\hat{x}_1^k$  respectively. Let us consider a *coupling* of the two conditional distributions  $P(\cdot | X_1^k = x_1^k)$  and  $P(\cdot | X_1^k = \hat{x}_1^k)$  and denote this by  $\pi_k^n(\cdot | x_1^k, \hat{x}_1^k)$ . This can be thought of as the joint distribution of a pair of random variable sequences,  $(Y_k^n, \hat{Y}_k^n)$ .

Now for  $k < i \leq n$ , define the “influence” coefficients  $v_{i,k}(x_1^k, \hat{x}_1^k)$  by:

$$v_{i,k}(x_1^k, \hat{x}_1^k) := \left( \sum_{y_k^n} \left( \pi_k^n[\hat{Y}_i \neq y_i | Y_k^n = y_k^n] \right)^2 P[Y_k^n = y_k^n | Y_1^k = x_1^k] \right)^{1/4} \quad (5)$$

Put  $v_{k,k}(x_1^k, \hat{x}_1^k) := 1$ . Set

$$v_{i,k} := \max_{x_1^k, \hat{x}_1^k} v_{i,k}(x_1^k, \hat{x}_1^k).$$

Now define

$$U := \max_i \sum_{1 \leq \ell \leq i} v_{i,\ell}^2, \quad (6)$$

and

$$V := \max_k \sum_{k \leq i \leq n} v_{i,k}^2. \quad (7)$$

Then the generalisation of (4) is [6, Theorem3.1]:

$$d_T(P, Q) \leq \sqrt{UV} \sqrt{2H(P | Q)}. \quad (8)$$

As an example of this inequality, consider the space of the  $m$ -fold product of the symmetric group  $S_n$ . The measure on each component space is the uniform measure and on the product space we take the product measure. To compute the coefficients (5), let us expose the variables component by component. Having exposed all variables in the first  $t < m$  components, suppose we have further exposed  $k < n$  variables in the  $t + 1$ st component. We take the natural coupling that simply swaps the permutation in this component and is the identity everywhere else. For this coupling, an easy calculation shows that

$$v_{i,k}(x_1^k, \hat{x}_1^k) = 1/\sqrt{n-k},$$

for  $k$  within this component and zero everywhere else. Thus

$$\begin{aligned} U &= \max_i \sum_{\ell < i} v_{i,\ell}^2 \\ &= \max_i \sum_{1 \leq \ell < i} \frac{1}{n-\ell} \\ &= H_n, \end{aligned}$$

where  $H_n \approx \log n$  is the  $n$ th Harmonic number and

$$\begin{aligned} V &= \max_k \sum_{i > k} v_{i,\ell}^2 \\ &= \max_k \sum_{i > k} \frac{1}{n-k} \\ &= 1. \end{aligned}$$

Hence on the product space  $S_n^m$  of the  $m$ -fold product of the symmetric group, we get the inequality:

$$d_T(P, Q) \leq \sqrt{\log n} \sqrt{2H(P | Q)} \quad (9)$$

Actually Talagrand is able to prove a stronger result on the symmetric group. For the  $m$ -fold product of the symmetric group, he proves:

$$P(A)P(B) \leq \exp\left(-d_T^2(A, B)/16\right) \quad (10)$$



There is a simple packaged form in which it is often convenient to apply these inequalities. In this form, the two components of the analysis are separated from each other:

- The *probabilistic* component which gives a concentration of measure inequality.
- The *smoothness* of the function of interest.

In applications, the concentration of measure inequality is taken ready-made without extra effort and one has only to verify the smoothness conditions on the function, which is often very easy.

**Theorem 1** *Let  $f$  be a real-valued function on a product space  $\Omega = \prod_{i \in [n]} \Omega_i$  with a measure  $P$  (which is not necessarily with product measure). Suppose for each  $x \in \Omega$ , there are reals  $\alpha_i(x), i \in [n]$  such that any **one** of the following conditions holds:*

$$f(x) \leq f(y) + \sum_{x_i \neq y_i} \alpha_i(x), \quad \text{for all } y \in \Omega, \quad (11)$$

or,

$$f(x) \geq f(y) - \sum_{x_i \neq y_i} \alpha_i(x), \quad \text{for all } y \in \Omega. \quad (12)$$

- *If a measure concentration inequality holds in the form*

$$P(A)P(B) \leq \exp\left(-d_T^2(A, B)/\alpha\right), \quad (13)$$

*for some  $\alpha > 0$ , and uniformly for all  $x \in \Omega$ ,*

$$\sum_i \alpha_i^2(x) \leq c, \quad (14)$$

*then we have the following concentration result for  $f$  around its median value  $\mathbf{M}[f]$ :*

$$P[|f - \mathbf{M}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{\alpha c}\right). \quad (15)$$

- If a measure concentration inequality holds in the form

$$d_T(P, Q) \leq \alpha \sqrt{2H(P \mid Q)} \quad (16)$$

for some  $\alpha > 0$  and also the coefficients  $\alpha_i(x)$  satisfy

$$\mathbf{E}[\sum_i \alpha_i^2(X)] \leq c, \quad (17)$$

then, we have the following concentration result for  $f$  around its mean:

$$P[|f - \mathbf{E}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{\alpha c}\right). \quad (18)$$

Note two features of the condition (11) or (12): first the asymmetry: we only need one of these one-sided versions to hold. Second its non-uniformity: the coefficients  $\alpha_i$  are allowed to depend on  $x$ . Both these features contribute to the power of the inequalities in applications.

## 4 High Probability Analyses

### 4.1 Top vertices

The analysis is particularly easy when  $v$  is a top vertex in Algorithm P. For, in this case, the incident edges all receive colours independently of each other. This is exactly the situation of the classical balls and bins experiment: the incident edges are the “balls” that are falling at random independently into the colours that represent the “bins”. Let  $T_e, e \in E$ , be the random variables taking values in  $[\Delta]$  that represent the tentative colours of the edges. Then the number of edges unsuccessfully coloured around  $v$  (and hence the new degree) is a function  $f(T_e, e \in N^1(v))$ , where  $N^1(v)$  denotes the set of edges incident on  $v$ . It is easily seen that this function has the *Lipschitz* property with constant 1: changing only one argument while leaving the others fixed only changes the value of  $f$  by at most 1. Thus:

- We can take all coefficients  $\alpha_i = 1$  in (11).
- Since the variables  $T_e, e \in N^1(v)$  are independent, we can apply the Talagrand inequality (2) for the product spaces  $[\Delta]^{N^1(v)}$  when  $v$  is a “top” vertex.

Hence, applying the first part of Theorem 1, we get the following sharp concentration result easily:

**Theorem 2** *Let  $v$  be a top vertex in algorithm P and let  $f$  be the number of edges around  $v$  that are successfully coloured in one round of the algorithm. Then,*

$$\Pr[|f - \mathbb{E}[f]| > t] \leq 2 \exp\left(\frac{-t^2}{4\Delta}\right),$$

For  $t := \epsilon\Delta$  ( $0 < \epsilon < 1$ ), this gives an exponentially decreasing probability for deviations around the mean. If  $\Delta \gg \log n$  then the probability that the new degree of any vertex deviates far from its expected value is inverse polynomial, i.e. the new max degree is sharply concentrated around its mean.

## 4.2 Bottom Vertices

The analysis for the “bottom” vertices in Algorithm P is more complicated in several respects. It is useful to see why so that one can appreciate the need for a more sophisticated analysis.

To start with, one could introduce an indicator random variable  $X_e$  for each edge  $e$  incident upon a bottom vertex  $v$ . These random variables are not independent however. Consider a four cycle with vertices  $v, a, w, b$ , where  $v$  and  $w$  are bottom vertices and  $a$  and  $b$  are top vertices. Let’s refer to the process of selecting the winner (step 2 of the algorithm P) as “the lottery”. Suppose that we are given the information that edge  $va$  got tentative colour red and lost the lottery—i.e.  $X_{va} = 0$ —and that edge  $vb$  got tentative colour green. We’ll argue intuitively that given this, it is more likely that  $X_{wb} = 0$ . Since edge  $va$  lost the lottery, the probability that edge  $wa$  gets tentative colour red increases. In turn, this increases the probability that edge  $wb$  gets tentative colour green, which implies that edge  $vb$  is more likely to lose the lottery. So, not only are the  $X_e$ ’s not independent, but the dependency among them is particularly malicious.

One could hope to bound this effect by using Talagrand’s inequality in its simplest form. This is also ruled out however, for two reasons. The first is that the tentative colour choices of the edges around a vertex are not independent. This is because the edges incident on a vertex are assigned a permutation of the colours. The second reason applies even if we pretend that all edges act independently. The new degree of  $v$ , a bottom vertex

in algorithm P is a function  $f = f(T_e, e \in N(v))$ , where  $N(v)$  is the set of edges at distance at most 2 from  $v$ . Thus  $f$  depends on as many as  $\Delta(\Delta - 1) = \Theta(\Delta^2)$  edges. Even if  $f$  is Lipschitz with constants  $d_i = 1$ , this is not enough to get a strong enough bound because  $d = \sum_i d_i^2 = \Theta(\Delta^2)$ . Applying Theorem 1 as above would give the bound

$$\Pr[|f - \mathbb{E}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{\Theta(\Delta^2)}\right).$$

This bound however is useless for  $t = \epsilon \mathbb{E}[f]$  since  $\mathbb{E}[f] \approx \Delta/e$ .

We will use Marton's extension of Talagrand's inequality to handle the dependence and we shall use the asymmetric and non-uniform nature of the condition (11) to control the effects of the individual random choices much more effectively.

Let  $N^1(v)$  denote the set of "direct" edges—i.e. the edges incident on  $v$ —and let  $N^2(v)$  denote the set of "indirect edges" that is, the edges incident on a neighbour of  $v$ . Let  $N(v) := N^1(v) \cup N^2(v)$ . The number of edges unsuccessfully coloured at vertex  $v$  is a function  $f(T_e, e \in N(v))$ .

For a tentative colouring  $T_e = c_e$ , choose the coefficients  $\alpha_i(c)$  as follows. Recall that edges compete in a "lottery" at the top vertices: for each colour, one "winner" is picked by a top vertex out of all the edges incident on it that receive the same tentative colour. Choose  $\alpha_i$  to be 1 for all unsuccessfully coloured edges around  $v$  and for all "winners"  $(w, z)$  such that the edge  $(v, w)$  took part in the lottery that  $(w, z)$  won (hence  $(w, z)$  was responsible for  $(v, w)$  being unsuccessful and serves as a *witness* for this fact). All other  $\alpha_i$  are 0. Thus at most  $2\Delta$  edges have non-zero  $\alpha_i$  values, and hence  $\sum_i \alpha_i^2 \leq 2\Delta$ . To see that (11) holds, let us look at an unsuccessfully coloured edge  $e$  in the colouring  $x$ . If it is also unsuccessful in the colouring  $y$ , it is counted in the term  $f(y)$ . Otherwise, at least one of  $e$  or its "witness"  $e'$  must be coloured differently in  $y$  and this will be counted in the second term in the right-hand side.

The underlying space is  $S_{\Delta}^{\bar{N}^2(v)}$  where  $\bar{N}^2(v) := \{u \mid d(u, v) = 2\} \cup \{v\}$ . Now applying the measure concentration result for the  $m$ -fold product of the symmetric group  $S_{\Delta}$  from Marton's inequality, namely (9) and using the second part of Theorem 1, we arrive at the following sharp concentration result:

**Theorem 3** *Let  $v$  be a bottom vertex in algorithm P and let  $f$  be the*

number of edges successfully coloured around  $v$  in one stage of either algorithm. Then,

$$\Pr[|f - \mathbb{E}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{2\Delta\sqrt{\log \Delta}}\right).$$

For  $t = \epsilon\Delta$ , this gives a probability that decreases almost exponentially in  $\Delta$ . As remarked earlier, if  $\Delta \gg \log n$ , this implies that the new max degree is sharply concentrated around the mean (with failure probability roughly inverse polynomial in  $n$ ).

One can improve this by using the stronger inequality (10) for the symmetric group. This gives the bound:

$$\Pr[|f - \mathbb{M}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{32\Delta}\right).$$

We first gave the result with Marton's inequality because although it gives a somewhat weaker bound in this example, it actually illustrates a general method with wider applicability.

## 5 A General Framework

The method used in the last section is actually applicable in a much more general way to the analysis of a wide range of randomised algorithms in a distributed setting. In this section, we outline a fairly general framework in which such concentration results can be directly inferred. Let  $f$  be a function to be computed by a randomised local algorithm in a distributed environment represented by a graph  $G = (V, E)$ . We shall lay down conditions on  $f$  and on algorithms computing  $f$  locally that will enable the methods in the previous section to be extended to derive a sharp concentration result on  $f$ . In particular, we indicate how the edge colouring algorithm above [8] as well as the edge and vertex colouring algorithms from [2, 3] follow directly as well as the analysis of a vertex colouring process in [7].

Suppose that  $f$  is a function determined by each vertex  $v$  of the graph assigning labels  $\ell(v)$  to itself and labels  $\ell(v, e)$ ,  $e \in N(v)$  to its incident edges by some randomised process. In the edge colouring problem, the labels on the edges are their colours (we may assume for instance that the lower numbered

vertex assigns the colour to an incident edge) and the vertex labels are empty; in the vertex colouring problems, the vertex labels are the colours and the edge labels are empty.

The two necessary and sufficient conditions for sharp concentration of  $f$  are as follows:

**L: Locality of the function:** The function  $f$  is *Lipschitz*: changing any one label changes the value of  $f$  by at most 1. Furthermore, following Spencer [9],  $f$  is *h-certifiable* for some function  $h : R \rightarrow R$  i.e. for each  $x$ , there is a subset  $I = I(x)$  (the “certificate”) of the labels of size at most  $h(f(x))$  such that for any other  $y$  agreeing with  $x$  on  $I$ ,  $f(y) \geq f(x)$ . In edge colouring, the function  $f$  is the number of edges unsuccessfully coloured around a vertex. It is clearly Lipschitz. Also it is  $h$  certifiable for  $h(x) = 2x$  since each unsuccessfully coloured vertex can be attributed to one other adjacent edge of the same colour.

**M: Concentration of Measure in the Probability Space:** There is concentration of measure in the underlying space in one of the two forms (13) or (16) with some positive coefficient  $\alpha$ . Marton’s inequality gives a method for determining this coefficient as  $\alpha = \sqrt{UV}$  for  $U$  and  $V$  determined by (5) through (7). Two particularly important cases in which this obtains are the following. First we assume that the vertex label assigned by a vertex is independent of the edge labels it assigns and that each vertex acts independently of the others.

**I:** The labels  $\ell(v, e)$  assigned by a vertex to its incident edges are independent of each other. In this case, the algorithm defines a fully independent probability space and Talagrand’s inequality (2) applies.

**S:** The labels  $\ell(v, e)$  assigned by a vertex to its incident edges are given by a permutation distribution. In this case, the algorithm is *symmetric* with respect to the labels. Marton’s theorem yields measure concentration in the form (16) with  $\alpha = \sqrt{\log n}$  and Talagrand’s inequality for the symmetric group gives measure concentration in the form (13) with  $\alpha = 16$ .

The condition **I** obtains in the algorithms in [3, 7] while the condition **S** obtains in the algorithm of [8] discussed in detail above.

**Theorem 4** *Let  $f$  be a Lipschitz function which is  $h$ -certifiable computed by an algorithm satisfying the measure concentration property  $\mathbf{M}$  for some  $\alpha > 0$ . Then*

$$\Pr[|f - \mathbf{M}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{\alpha h(\mathbf{M}[f])}\right).$$

*In particular if the algorithm satisfies either full independence  $\mathbf{I}$  or is symmetric, satisfying  $\mathbf{S}$ , then we have the concentration result:*

$$\Pr[|f - \mathbf{M}[f]| > t] \leq 2 \exp\left(-\frac{t^2}{ch(\mathbf{M}[f])}\right),$$

*where the constant  $c$  is 4 for the independent case and 16 for the symmetric case.*

## 6 Acknowledgement

I would like to thank Kati Marton for her invaluable help in understanding how to apply her result. I would like to thank Alessandro Panconesi for his help in drafting this paper, for invaluable discussions and for resolutely prodding me to pursue the “demise of the permanent” from [8]. Now, finally, R.I.P.

## References

- [1] Bollobas, B. : Graph Theory: An Introductory Course. Springer-Verlag 1980.
- [2] Dubhashi, D., Grable, D.A., and Panconesi, A.: Near optimal distributed edge colouring via the nibble method. Theoretical Computer Science 203 (1998), 225–251, a special issue for ESA 95, the 3rd European Symposium on Algorithms.
- [3] Grable, D., and Panconesi, A: Near optimal distributed edge colouring in  $O(\log \log n)$  rounds. Random Structures and Algorithms **10**, Nr. 3 (1997) 385–405

- [4] Luby, M.: Removing randomness in parallel without a processor penalty. *J. Computer and Systems Sciences* **47:2** (1993) 250–286.
- [5] Marton, K.: Bounding  $\bar{d}$  distance by informational divergence: A method to prove measure concentration. *Annals of Probability* **24** (1996) 857–866.
- [6] Marton, K.: On the measure concentration inequality of Talagrand for dependent random variables. submitted for publication. (1998).
- [7] Molloy, M. and Reed, B.: A bound on the strong chromatic index of a graph. *J. Comb. Theory (B)* **69** (1997) 103–109.
- [8] Panconesi, A. and Srinivasan, A.: Randomized distributed edge coloring via an extension of the Chernoff–Hoeffding bounds. *SIAM J. Computing* **26:2** (1997) 350–368.
- [9] Spencer, J.: Probabilistic methods in combinatorics. In proceedings of the International Congress of Mathematicians, Zurich, Birkhauser (1995).
- [10] Talagrand, M.: Concentration of measure and isoperimetric inequalities in product spaces. *Publ. math. IHES*, **81:2** (1995) 73–205..



## Recent BRICS Report Series Publications

- RS-98-24 Devdatt P. Dubhashi. *Talagrand's Inequality and Locality in Distributed Computing*. October 1998. 14 pp.
- RS-98-23 Devdatt P. Dubhashi. *Martingales and Locality in Distributed Computing*. October 1998. 19 pp.
- RS-98-22 Gian Luca Cattani, John Power, and Glynn Winskel. *A Categorical Axiomatics for Bisimulation*. September 1998. ii+21 pp. Appears in Sangiorgi and de Simone, editors, *Concurrency Theory: 9th International Conference, CONCUR '98 Proceedings*, LNCS 1466, 1998, pages 581–596.
- RS-98-21 John Power, Gian Luca Cattani, and Glynn Winskel. *A Representation Result for Free Cocompletions*. September 1998. 16 pp.
- RS-98-20 Søren Riis and Meera Sitharam. *Uniformly Generated Submodules of Permutation Modules*. September 1998. 35 pp.
- RS-98-19 Søren Riis and Meera Sitharam. *Generating Hard Tautologies Using Predicate Logic and the Symmetric Group*. September 1998. 13 pp.
- RS-98-18 Ulrich Kohlenbach. *Things that can and things that can't be done in PRA*. September 1998. 24 pp.
- RS-98-17 Roberto Bruni, José Meseguer, Ugo Montanari, and Vladimiro Sassone. *A Comparison of Petri Net Semantics under the Collective Token Philosophy*. September 1998. 20 pp. To appear in *4th Asian Computing Science Conference, ASIAN '98 Proceedings*, LNCS, 1998.
- RS-98-16 Stephen Alstrup, Thore Husfeldt, and Theis Rauhe. *Marked Ancestor Problems*. September 1998.
- RS-98-15 Jung-taek Kim, Kwangkeun Yi, and Olivier Danvy. *Assessing the Overhead of ML Exceptions by Selective CPS Transformation*. September 1998. 31 pp. To appear in the proceedings of the *1998 ACM SIGPLAN Workshop on ML*, Baltimore, Maryland, September 26, 1998.
- RS-98-14 Sandeep Sen. *The Hardness of Speeding-up Knapsack*. August 1998. 6 pp.