Branching random walk and searching in trees: final report

L. Addario-Berry (McGill University),
N. Broutin (INRIA Rocquencourt)
L. Devroye (McGill University)
C. McDiarmid (University of Oxford)

September 22, 2010

1 Overview of the subject

A branching random walk is a Galton-Watson tree T to which the individuals have been assigned spatial positions (in $\mathbb{R}$, say), in the following manner. The root r is placed at the origin. Each child c of the root is independently given a random position P_c; the distribution of each such displacement is given by some real random variable X. More generally, when an individual u has a child v, v appears at position $P_v = P_u + X_v$, where X_v is an independent copy of X. This yields a natural, though idealized, discrete model of how a population may diffuse over time.

Branching random walks are a natural and basic object of study in probability, and are far from being fully understood. Furthermore, branching random walks turn out to have strong connections in other parts of mathematics and theoretical computer science. To highlight a particularly notable example, consider the problem of understanding the *minimum (most negative) position* of any individual in the n'th generation of T, which we denote M_n. An understanding of this random variable ends up being fundamental for analyzing the expected worst-case behavior of a host of data structures of great interest to the theoretical computer science community. The behavior of the expected value $\mathbb{E}M_n$ also turns out to be intimately connected to the uniqueness of "travelling-wave" solutions to a reaction-diffusion equation called the Kolmogorov–Petrovskii–Piskounov equation, given by $\frac{\partial u}{\partial t} = \frac{1}{2}\frac{\partial^2 u}{\partial x^2} + f(u)$, taking solutions $u(x,t) : \mathbb{R} \times \mathbb{R}^+ \to [0,1]$. This equation has received much attention using both probabilistic and analytical techniques.

The second part of the workshop title refers to the closely related problem of *finding* individuals, or paths, in a tree T, that have particular, desirable properties. Though we may know via probabilistic arguments that with very high probability, M_n is at least x, it may be impractical to actually find an individual with displacement x. For example, if we are searching in a branching random walk tree whose branching distribution B satisfies $\mathbb{E}B > 1$, then the expected number of individuals in generation n grows exponentially in n. If there is only one v in generation n with displacement $P_v \leq x$, then to find this individual may take exponential time. There are several known approaches for finding individuals that come close to attaining the minimal displacement while using limited resources; however, a general understanding of optimal strategies for finding good approximations has yet to be developed.

Very recently, Bramson and Zeitouni (2007,2008+), Addario-Berry and Reed (2009), and Hu and Shi (2008+) have all proved related results on the position of the minimum M_n and on its concentration around its mean; these results build on existing work by Biggins, Devroye, and McDiarmid, among others. It is now clear that for a wide range of branching random walks, there are constants $\alpha \in \mathbb{R}$ and $\beta > 0$ such that M_n is roughly $\alpha n + \beta \log n$ in expectation and in probability (though, as established by Hu and Shi, *not* almost surely). Furthermore, $\mathbf{P}\{M_n - \mathbf{E}M_n > x\}$ decays exponentially quickly in x.

These recent results are established using very different techniques. Bramson and Zeitouni prove their concentration result by connecting the BRW to a family of recursive equations and then studying a related Lyapunov function; Addario-Berry and Reed use a rather combinatorial argument based on the second moment method; and Hu and Shi proceed by connecting a certain exponential martingale with a distinguished path called the "spine" of the tree. There are some intriguing links between the three methods, and a deeper understanding of these links has the potential to lead to a common strengthening of all three results.

From the algorithmic point of view, it is natural to ask whether we can efficiently *find* an individual in the n'th generation who attains or comes close to attaining the minimum position M_n. One extremely natural technique for doing finding an individual with small displacement is the following: given all the individuals seen thus far, consider the individual with the minimum displacement and explore its children. We refer to this approach as the *greedy algorithm*. Aldous (1992) first analyzed this procedure, in the context of binary trees, and gave sufficient conditions for the greedy algorithm to find a path approximating the path to M_n up to a fixed error ϵn in the slope. Aldous further made the (very natural) conjecture that the greedy algorithm stochastically optimizes the minimum position seen up to time t. Surprisingly, this conjecture turns out to be false, as was shown by Stacey (1999); however, Stacey also showed that the greedy algorithm is optimal in a much weaker sense than that proposed by Aldous, and posed multiple intriguing questions about the strongest sense in which some form of optimality holds. (It is worth mentioning the work of Karp and Pearl (1983) and McDiarmid (1992), who studied the behavior of "bounded backtracking" algorithms for finding individuals with small displacement, also using the connection with branching random walks.)

More recently, Pemantle (2008+) has proposed a new algorithm, also of a "greedy" variety for finding individuals whose displacement is close to M_n; the algorithm applies in the case that the displacements are Bernoulli. Pemantle has proved (for certain "subcritical" choices of the Bernoulli parameter) that the running time of this algorithm is with high probability best possible up to $o(n)$ terms. A key step that will be required for the analysis of the remaining cases of this algorithm is to understand branching random walks in which all particles whose displacement exceeds a fixed slope are "killed" and no longer reproduce. Pemantle has also made an intriguing conjecture about the *class* of algorithms in which optimal algorithms must lie.

When using results on branching random walk to analyze (usually tree-based) data structures, one of the primary obstacles is the fact that branching random walks are usually limit objects for the data structures in question. To apply results for BRW to the data structures themselves, one must somehow "pass from infinite to finite", and it is usually far from clear how to proceed. For example, the expected worst-case query time is essentially determined by the finitization of the minimum displacement M_n. Using this fact, it is rather straightforward to bound the first-order asymptotics of the expected worst-case query time (i.e. to "finitize" the fact that $M_n = (1 + o(1))\alpha n$), but understanding lower order terms can be extremely difficult. For the case of binary search trees, the link with (binary branching, exponential displacement) branching random walks has been known since work of Biggins, Devroye, and Pittel in the 1980's. However, pinning down the lower-order behavior of worst-case query time was only accomplished (by Reed and, later, Drmota) in 2003. A substantial part of the technical challenge in Reed's approach was dealing with the finitization, and his solution made heavy use of binary branching. Chauvin and Drmota (have had some success at extending Drmota's approach to more general search trees; however, a general strategy for handling this finitization remains elusive.

Finally, one may consider a generalization of the branching random walk model, where the displacements are given by a Markov chain (displacements along distinct edges remain independent), or where the random displacements depend on the initial position. In this case much less is known about the behavior of the system. The first-order asymptotics can be derived without much more difficulty than in branching random walks, but lower order behavior seems much harder to handle. A common approach is to renormalize so that the first order term vanishes (i.e. so that $M_n = o(n)$) by subtracting a fixed constant from each displacement. Once this is done, even the question of recurrence or transience is rather difficult; though at this point it is rather well understood due to work of Comets, Menshikov and Popov (1998), Machado and Popov (2000,2001,2003), Gantert and Müller (2006) and Müller (2008+).

2 Overview of the workshop

During the week of the workshop there was a strong emphasis on direct contact and collaboration; as such, we limited talks to the mornings of the workshop, and spent the afternoons discussing and working in groups. For the first two afternoons we had problem sessions, to help motivate future discussions. What follows is a summary of some of the substantial open problems that were presented during the workshop.

3 Overview of some of the open problems arising from or presented at the workshop.

ARCSINE TYPE LAW IN A BRANCHING RANDOM WALK?
Yueyun Hu

Let $(V(x),\, x \in \mathbb{T})$ be a branching random walk on the real line. The positions of the particles in the n-th generation are denoted by $(V(x),\, |x| = n)$. Let $[\![\emptyset,\, x]\!] = \{x_0, x_1, ..., x_n\}$ be the set of vertices on the shortest path connecting the origin $\emptyset$ to x such that $|x_i| = i$ for $i = 0, ..., n$. Define $\overline{V}(x) := \max_{y\in[\![\emptyset,\, x]\!]} V(y)$, for $x \in \mathbb{T}$. It has been shown in Hu and Shi (2007, Ann. Prob.) that $\min_{|x|=n} \overline{V}(x)$ plays an important role in the study of random walk in random environment on the tree $\mathbb{T}$. Recently, Fang and Zeitouni (2009, Arxiv) and Faraud, Hu and Shi (2010, preprint) have independently obtained the asymptotic behaviors of the min-max $\min_{|x|=n} \overline{V}(x)$. Let $x^{*,n}$ be a vertex in nth generation satisfying $\overline{V}(x^{*,n}) = \min_{|x|=n} \overline{V}(x)$. We are interested in

$$\tau_n := \min\{j \le n : V(x_j^{*,n}) = \overline{V}(x^{*,n})\}.$$

Question: Does $\frac{\tau_n}{n}$ converge as $n \to \infty$? If yes, what is the limiting distribution?

RECONSTRUCTING SCENERY FROM A MARKED GENEALOGICAL TREE
Serguei Popov

Let us put 0s and 1s to the sites of $\mathbb{Z}^d$, $d \ge 2$, at random (e.g., independently and with equal probabilities). This "coloring" of $\mathbb{Z}^d$ is referred to as *scenery*. Then, consider a branching random walk starting from one particle at the origin: the transition probabilities are those of the simple random walk, and each particle is substituted by two on each step. Then, to each vertex of the genealogical tree of the BRW (which is a rooted binary tree in this case), we put 0 or 1 according to the color the corresponding particle observes. The problem is: having only this genealogical tree with marks, obtain an algoritheorem that a.s. reconstructs the original scenery (up to shift/reflection).

FROGS
Nina Gantert

The frog model can be described as follows. Let G be a graph and take one vertex to be the origin. Initially there is a number of sleeping particles ("frogs") at each site of the graph G except at the origin. The origin contains one active frog. The active frog then starts a discrete-time simple random walk on the vertices of G. Each time an active frog visits a site with sleeping frogs the latter become active and start moving according to the same random walk as the active frogs, independently from everything else. An interpretation of the model is the distribution of information: Active frogs hold some information and share them with sleeping frogs as soon as they meet. The sleeping frogs become active and start helping in the process of spreading the information (cf. [21]). The frog model can also be interpreted as a "once-branching" random walk, i.e. a branching random walk where branching takes only place in a site which is visited for the first time.
We denote by η_x the number of sleeping frogs initially in x. One is interested in recurrence and transience, i.e. whether the probability of having infinitely many visits to the origin is 1 or strictly less than 1.

For a symmetric random walk on $\mathbb{Z}^d$, the frog model (starting with one frog at each site) is known to be recurrent, cf. [19] and [20]. There are variants of the model where the frogs have random lifetimes, and one is interested in survival/extinction of the process (and its dependence on the parameters), see [21]. Another question which has been investigated for the model is the existence of shape theorems.

Open problems

(1) Consider the frog model with simple random walk. Not even for transitive graphs, starting with one frog everywhere, recurrence and transience are settled: an example of a graph where this question is open is the binary tree.
(2) A natural conjecture is the following: Assume that the graph G is transitive, the underlying random walk is homogeneous and the initial configuration η is i.i.d. Then we have either

$$P_\eta\left[\text{the origin is visited infinitely often}\right] = 0 \qquad \mu\text{-a.s.}$$

or

$$P_\eta\left[\text{the origin is visited infinitely often}\right] = 1 \qquad \mu\text{-a.s.}$$

For background on the model and further open problems, we refer to [21] and [18].

SURVIVAL OF BRANCHING RANDOM WALKS ON GRAPHS
Sebastien Mueller

A branching random walk (BRW) on a graph G is defined as follows: start the process with one particle in the origin, then inductively each particle dies at rate 1 and gives birth to new particles at each neighboring vertex at rate λ independently of the rest of the process. There exists some $\lambda_c = \lambda_c(G)$ such that the process dies out a.s. if $\lambda < \lambda_c$ and survives with positive probability if $\lambda > \lambda_c$.

Open question 1: What is λ_c?

One may start with the following:

Open question 1a: Let the underlying graph be a realization $\mathcal{T}$ of a Galton–Watson tree. What is $\lambda_c(\mathcal{T})$?

Open question 2: What happens in the critical case, i.e., $\lambda = \lambda_c$?
Background
The branching random walk (BRW) on a locally finite graph $G = (V, E)$ is a continuous-time Markov process whose state space is a suitable subset of $\mathbb{N}^V$. The BRW with parameter λ is described through the number $\eta(t, v)$ of particles at vertex v at time t and evolves according the following rules: for each $v \in V$

$$\begin{aligned}
\eta(t,v) &\rightarrow \eta(t,v) - 1 \text{ at rate } \eta(t,v)\\
\eta(t,v) &\rightarrow \eta(t,v) + 1 \text{ at rate } \lambda \sum_{u:uv\in E} \eta(t,u).
\end{aligned}$$

Let $o \in V$ be some distinguished vertex of G and denote $\mathbb{P}_o$ the probability measure of the process started with one particle at o at time 0. Define $Z(t) = \sum_v \eta(t, v)$ the number of particles at time t. One says the BRW survives (with positive probability) if

$$\mathbb{P}_o\left(Z(t) > 0 \text{ for all } t \geq 0\right) > 0.$$

Define the *first moment matrix* $M = (m(x, y))_{x,y\in V}$: $m(u, v) := \lambda$ if $uv \in E$ and $m(u, v) := 0$ otherwise. In other words: $M = \lambda \cdot A$, where A is the adjacency matrix of the graph $G = (V, E)$. Survival of a branching processes is connected to the (expected) growth rate. It turns out that the following definition is very useful:

$$\lambda(M) := \inf\{\lambda > 0 : \exists 0 < f \in L_\infty : \ \lambda M f \geq f\}$$

Bertacchi and Zucca [22] answered Question 1 in the following way:

Theorem 1. $\lambda_c = \lambda(M)$

Despite the latter theorem there are no known *non-trivial* examples where one can really calculate or describe explicitly the critical value λ_c. Observe that for all Cayley graphs, transitive graphs or even quasi-transitive graphs one does have a good description of λ_c, see [22] or earlier work [25], [23]. Furthermore, for percolation clusters on graphs with bounded degree it turns out that $\lambda_c = 1$.

A natural candidate for a non-trivial example is the BRW on a Galton–Watson tree. This model was studied in Pemantle and Stacey [24]. While they obtained a partial description of λ_c, Question 1a is open in the sense that no general formula for λ_c in terms of the offspring distribution of the Galton–Watson process is known.

The method used in [22] does not seem to carry over to the critical case. In [22], they also give an example (in a more general setting) that survival may also occur in the critical case. This example does not apply for graphs with bounded degrees, so that one might conjecture that in our setting the BRW always dies out in the critical case.

Searching in random trees
COLIN MCDIARMID

Consider an infinite complete binary tree, with iid edge-lengths X_e. For each node v let its 'position' or 'birth time' $S(v)$ be the sum of the edge-lengths along the path to it from the root. Let M_n denote the first birth time of a node in generation n (at depth n). How quickly can we find a node v in generation n with $S(v)$ equal to M_n or nearly so?

The setting is a little too special for some algoritheoremic questions since there are no dead-ends, so let us generalise to a Galton-Watson tree with bounded family size and mean family size $m > 1$, conditioned on survival.

Suppose that the edge-lengths are Bernouilli. The problem is easy if $m\,\mathbb{P}(X_e = 0) > 1$, as the first birth times M_n stay bounded, and it is easy to find an optimal node; and further, in the case $m\,\mathbb{P}(X_e = 0) = 1$ the first birth times grow only like $\log\log n$ and again it is easy to find an optimal node, see [31, 32].

Let us suppose then that $m\,\mathbb{P}(X_e = 0) < 1$. It has been known since the work of Hammersley, Kingman and Biggins in the mid 70s that, conditional on survival, there is a constant $\gamma > 0$ such that $M_n/n \to \gamma$ a.s. (and we may specify γ). Further, for any $\epsilon > 0$ there is a polynomial time algoritheorem, which, with high probability conditional on survival, finds a node v in generation n with $S(v) < (1+\epsilon)\gamma n$, see [31, 32]. Here 'time' refers to number of edge-lengths examined.

This seemed a satisfactory algoritheoremic result: the optimal value was known up to an error term $o(n)$, and we could quickly find a node with birth time at most a similar distance from the optimum. But now we know the optimum value very precisely. Following limited earlier steps in [33], it is shown by Addario-Berry and Reed in [26] that M_n has expected value $\gamma n + \alpha \log n + O(1)$ for a given constant α, and M_n is strongly concentrated around this value.

Thus now it seems natural to seek a node v with $S(v)$ closer to the optimum value, and not settle for a birth time which is ϵn too big. There have been recent results suggesting that we can restrict a search to 'small' parts of the tree and still come close to an optimal solution, see [30, 29]. Also, see [27, 28, 35] for related work with a different algoritheoremic target (what is the best node you cacn find at any depth within n steps?)

Thus we ask how close to the optimum we can come in polynomial time? In other words, what is the smallest 'accuracy term' $\delta(n)$ for which there is a polynomial time algoritheorem which, with high probability conditional on survival, finds a node v in generation n with $S(v) \leq B(n) + \delta(n)$? The 'staged' back-tracking algoritheorem from [31, 32] will handle $\delta(n) = n \log\log n / \log n$, but what about say $\delta(n) = n^{\frac{1}{2}}$? Also, we could also abandon polynomial time, and ask say whether in time $O(2^{n^{\frac{1}{3}}})$ we can come within $n^{\frac{1}{3}}$ of the optimum value?

BRANCHING POINT PROCESSES
Louigi Addario-Berry

Let $\mathcal{U}$ be the *Ulam–Harris tree*, which has vertex set $\bigcup_{n=0}^{\infty} \mathbb{N}^n$ (we take $\mathbb{N}^0 = \emptyset$ by convention), is rooted at $\emptyset$, and in which a node $u_1 u_2 \ldots u_k$ has parent $u_1 u_2 \ldots u_{k-1}$ and children $\{u_1 u_2 \ldots u_k v, v \in \mathbb{N}\}$.

Let P be a point process on $\mathbb{R}$. We assume that P is such that almost surely, P has a (random) least element $p_0 > -\infty$ and that P almost surely has no accumulation points, so that it is possible to list the elements of P in increasing order as $\{p_i\}_{i\in\mathbb{N}}$. (If P has only finitely many elements, then we let $p_i = +\infty$ for all $i > |P|$.) We may then form a branching random walk with reproduction-diffusion mechanism P by associating to each node $v \in \mathcal{U}$ an independent copy P^v of $\mathcal{P}$. For a child vi of node v, we then view

p_i^v as the displacement from v to vi. For a node $v = u_1u_2 \dots u_k$, writing $v^0 = \emptyset$ and $v^i = u_1u_2 \dots u_i$ for $1 \le i \le k$, the *position* S_v of v is then $\sum_{i=0}^{k-1} p_{u_{i+1}}^{v^i}$, the sum of the displacements along the path from the root to v. Let M_n be the minimum position of any node in $\mathbb{N}^n$. By our assumptions on the point process, there is almost surely a node achieving this minimum value.

In a sequence of papers, each building on the results of the last, [36], [37] and [38] showed that under suitable conditions on the point process P, there is a finite constant γ such that, conditional on the survival of the branching process,

$$M_n/n \to \gamma \qquad \text{almost surely.}$$

When Hammersley initiated this research into the first-order behavior of M_n, he posed several questions to which complete answers remain unknown. In particular, he asked when more detailed information about $M_n - \gamma n$ than that given by the above law of large numbers can be found, about the expectation of M_n, and about whether the higher centralized moments of M_n are bounded. Thanks to the work of several researchers the answer to this question is now understood in great detail and in quite some generality *under the assumption that P is almost surely finite.* My question is about how easily this assumption can be removed.

Open question 1: Under what conditions on the point process P does $\mathbb{E}\{M_n\}$ have the form $an + b\log n + O(1)$, for some constants $a, b > 0$.

A technical difficulty to extending at least some of the techniques that work when P is almost surely finite, is the following. When P is almost surely finite, by randomly permuting the finite-displacement children of each node, it is possible to assume that in fact the finite displacements to the children of a node are identically distributed. This yields that the sequence of displacements down any non-backtracking path from the root using only finite-displacement edges, has precisely the behaviour of a random walk. Results on sample paths of random walks (a now-standard technique in analyzing branching random walks) can then be brought to bear on the problem. It turns out [**?**] that this technique can also be made to work when P is infinite if the inter-point spacings of P are independent and identically distributed. However, this only describes a relatively limited number of point processes. A potentially more powerful approach would be to extend existing results on sample paths of random walks to random-walk-like sequences of jumps where the jumps are not necessarily iid. As this is potentially a bit vague, here is a concrete open problem which could guide an attack via this sort of approach.

Open question 2: Let $(x_1, \dots, x_n)$ be real numbers with $\sum_{i=1}^n x_i = s \le \sqrt{n}$. Find sufficient conditions on $(x_1, \dots, x_n)$ to ensure that if $\sigma : [n] \to [n]$ is a uniformly random permutation, then

$$\mathbb{P}\left\{ \sum_{j=1}^{i} x_j > 0,\ 0 < j < n \right\} = \Omega\left(\frac{k}{n}\right).$$

References

[1] DJ Aldous. Greedy search on the binary tree with random edge-weights. *Combinatorics, Probability and Computing*, 1:281–293, 1992.

[2] Maury Bramson and Ofer Zeitouni. Tightness for the minimal displacement of branching random walk. *Journal of Statistical Mechanics: Theory and Experiment*, July 2007, 2007.

[3] Maury Bramson and Ofer Zeitouni. Tightness for a family of recursion equations. *Annals of Probability* 37(2):615-653, 2009.

[4] B. Chauvin and M. Drmota. The random multisection problem, travelling waves, and the distribution of the height of m-ary search trees. to appear in Algorithmica, 2005.

[5] F. Comets, MV Menshikov, and S.Y. Popov. One-dimensional branching random walk in random environment: a classification. *Markov Processes Relat. Fields*, 4(4):465–477, 1998.

[6] M. Drmota. An analytic approach to the height of binary search trees II. *Journal of the ACM*, 50: 333–374, 2003.

[7] N. Gantert and S. Muller. The critical branching Markov chain is transient. *Markov Processes and Related Fields*, 12:805–814, 2006.

[8] Yueyun Hu and Zhan Shi. Minimal position and critical martingale convergence in branching random walks, and directed polymers on disordered trees. *Annals of Probability* 37(2):742–789, 2009.

[9] H.K. Hwang. Second phase changes in random m-ary search trees and generalized quicksort: rates of convergence. *Annals of Probability*, 31(2):609–629, 2003.

[10] FP Machado and S.Y. Popov. One-dimensional branching random walks in a Markovian random environment. *Journal of Applied Probability*, 37(4):1157–1163, 2000.

[11] FP Machado, MV Menshikov, and S.Y. Popov. Recurrence and transience of multitype branching random walks. *Stochastic processes and their applications*, 91(1):21–37, 2001.

[12] FP Machado and S.Y. Popov. Branching random walk in random environment on trees. *Stochastic processes and their applications*, 106(1):95–106, 2003.

[13] CJH McDiarmid. Probabilistic analysis of tree search. *Disorder in Physical Systems*, pages 249–260, 1990.

[14] Sebastian Müller. Recurrence and transience for Branching Random Walks in an iid random environment. *Markov Processes and Related Fields*, 14:815–830, 2008.

[15] Robin Pemantle. Search cost for a nearly optimal path in a binary tree. submitted, 2008+.

[16] B. Reed. The height of a random binary search tree. *Journal of the ACM*, 50:306–332, 2003.

[17] A.M. Stacey. Searches on a binary tree with random edge-weights. *Combinatorics, Probability and Computing*, 8(06):555–565, 1999.

[18] N.GANTERT, P. SCHMIDT (2009) Recurrence for the frog model with drift on $\mathcal{Z}$. *Markov Processes and Related Fields 15 (1)*, 51-58.

[19] A. TELCS, N. WORMALD (1999) Branching and Tree-Indexed Random Walks on Fractals, *Journal of Applied Probability 36 (4)*, 999-1011

[20] S. POPOV (2001), Frogs in random environment, *Journal of Statistical Physics 102 (1/2)*, 191-201

[21] S. POPOV (2003), Frogs and some other interacting random walks models, *Discrete Mathematics and Theoretical Computer Science Proceedings AC*, 277-288

[22] D. Bertacchi and F. Zucca, Critical behaviors and critical values of branching random walks on multigraphs. *Journal of Applied Probability* 45, 2008.

[23] S. Müller. A criterion for transience of multidimensional branching random walk in random environment. *Electronic Journal of Probability*, 13:1189–1202, 2008.

[24] R. Pemantle and A. Stacey. The branching random walk and contact process on Galton-Watson and nonhomogeneous trees. *Annals of Probability.*, 29(4):1563–1590, 2001.

[25] A. Stacey. Branching random walks on quasi-transitive graphs. *Combin. Probab. Comput.*, 12(3):345–358, 2003. Combinatorics, probability and computing (Oberwolfach, 2001).

[26] Louigi Addario-Berry, Bruce Reed, Minima in Branching Random Walks, *Annals of Probability*, Volume 37, Issue 3 (2009), 1044-1079.

[27] D. Aldous, Greedy search on the binary tree with random edge-weights, *Combinatorics, Probability and Computing* **1** (1992) 281 – 293.

[28] D. Aldous, A Metropolis-type optimization algoritheorem on the infinite tree, *Algoritheoremica* **22** (1998) 388 - 412.

[29] Y. Hu and Z. Shi, Minimal position and critical martingale convergence in branching random walks, and directed polymers on disordered trees, arxiv:math/0702799.

[30] N. Gantert, Y. Hu and Z. Shi, Asymptotics for the survival probability in a killed branching random walk, arXiv:0811.0262v3.

[31] R.M. Karp and J. Pearl, Searching for an optimal path in a tree with random costs, *Artificial Intelligence* **21** (1983) 99 - 116.

[32] C. McDiarmid, Probabilistic analysis of tree search, in G. R. Grimmett and D.J.A. Welsh, editors, *Disorder in Physical Systems*, Clarendon Press, Oxford, 1990, 249 - 260.

[33] C. McDiarmid, Minimal positions in a branching random walk, *Ann. Appl. Prob.* **5** (1995), 128 – 139.

[34] C. McDiarmid and G. Provan, An expected cost analysis of backtracking and non-backtracking algoritheorems, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)* (1991), 172 – 177.

[35] R. Pemantle Search cost for a nearly optimal path in a binary tree, *Ann. Appl. Prob.* **19** (2009) 1273 –1291.

[36] J.M. Hammersley. Postulates for subadditive processes. *The Annals of Probability*, 2(4):652–680, 1974.

[37] J.F.C. Kingman. The first birth problem for an age-dependent branching process. *The Annals of Probability*, 3(5):790–801, 1975.

[38] J. D. Biggins. The first- and last-birth problems for a multitype age-dependent branching process. *Advances in Applied Probability*, 8(3):446–459, 1976.