

Towards a Theory of Random Walk Planning: Regress Factors, Fair Homogeneous Graphs, and Extensions

Hootan Nakhost^a and Martin Müller^a

^a *Department of Computing Science
University of Alberta, Edmonton, Canada
E-mails: nakhost@ualberta.ca,
mueller@ualberta.ca*

Random walks are a relatively new component used in several state of the art satisficing planners. Empirical results have been mixed: while the approach clearly outperforms more systematic search methods such as weighted A* on many planning domains, it fails in many others. So far, the explanations for these empirical results have been somewhat ad hoc. This paper proposes a formal framework for comparing the performance of random walk and systematic search methods. *Fair homogenous* and *Infinitely Regressable homogenous* graphs are proposed as graph classes that represents characteristics of the state space of prototypical planning domains, and is simple enough to allow a theoretical analysis of the performance of both random walk and systematic search algorithms. This gives well-founded insights into the relative strength and weaknesses of these approaches. The close relation of the models to some well-known planning domains is shown through simplified but semi-realistic planning domains that fulfill the constraints of the models.

One main result is that in contrast to systematic search methods, for which the branching factor plays a decisive role, the performance of random walk methods is determined to a large degree by the Regress Factor, the ratio between the probabilities of progressing towards and regressing away from a goal with an action. The performance of random walk and systematic search methods can be compared by considering both branching and regress factors of a state space.

1. Random Walks in Planning

Random walks, which are paths through a search space that follow successive randomized

state transitions, are a main building block of prominent search algorithms such as Stochastic Local Search techniques for SAT [1,2] and Monte Carlo Tree Search in game playing and puzzle solving [3,4,5,6].

Inspired by these methods, several recent satisficing planners also utilize random walk (RW) techniques. Identidem [7] performs a hill climbing search that uses random walks to escape from plateaus or saddle points. All visited states are evaluated using a heuristic function. Random walks are biased towards states with lower heuristic value. Roamer [8] enhances its best-first search (BFS) with random walks, aiming to escape from *search plateaus* where the heuristic is uninformative.

Arvand [9] takes a more radical approach: it relies exclusively on a set of random walks to determine the next state in its local search. For efficiency, it only evaluates the endpoints of those random walks. Arvand also learns to bias its random walks towards more promising actions over time, by using the techniques of *Monte Carlo Deadlock Avoidance* (MDA) and *Monte Carlo with Helpful Actions* (MHA). In the Arvand-RC system [10], local search is enhanced by the technique of *Smart Restarts*, and applied to solving Resource Constrained Planning (RCP) problems. *Arvand-LS* is a hybrid system which combines random walks with a local greedy best first search [11].

Compared to all other tested planners, Arvand-RC performs much better in RCP problems [10], which test the ability of planners in dealing with scarce resources. In IPC domains, RW-based planners tend to excel on domains with many paths to the goal. Scaling studies in [11] show that RW planners can solve much larger problem instances than other state of the art planners in the domains of *Transport*, *Elevators*, *Openstacks*, and *Visitall*. However, these planners perform poorly

in *Sokoban*, *Parking*, and *Barman*, puzzles with a small solution density in the search space.

While the success of RW methods in related research areas such as SAT and Monte Carlo Tree Search serves as a good general motivation for trying them in planning, it does not provide an explanation for why RW planners perform well. Previous work has highlighted three main advantages of random walks for planning:

- Random walks are more effective than systematic search approaches for escaping from regions where heuristics provide no guidance [7,8,9].
- Increased sampling of the search space by random walks adds a beneficial *exploration* component to balance the *exploitation* of the heuristic in planners [9].
- Combined with proper *restarting* mechanisms, random walks can avoid most of the time wasted by systematic search in dead ends. Through restarts, random walks can rapidly back out of unpromising search regions [7,9].

These explanations are intuitively appealing, and give a qualitative explanation for the observed behavior on planning benchmarks such as IPC and IPC-2011-LARGE [11]. Typically, random walk planners are evaluated by measuring their coverage, runtime, or plan quality in such benchmarks.

1.1. Studying Random Walk Methods

There are many feasible approaches for gaining a deeper understanding of these methods.

- Scaling studies, as in Xie et al. [11].
- Algorithms combining RW with other search methods, as in [8,12].
- Experiments on small finite instances where it is possible to “measure everything” and compare the choices made by different search algorithms.
- Direct measurements of the benefits of RW, such as faster escape from plateaus of the heuristic.
- A theoretical study of how RW and other search algorithms behave on idealized classes of planning problems which are amenable to such analysis.

The current paper pursues the latter approach. The main goal is a careful theoretical investigation of the first advantage claimed above - the question of how RW manage to escape from plateaus faster than other planning algorithms.

1.2. A First Motivating Example

As an example, consider the following well-known plateau for the FF heuristic, h_{FF} , discussed in [13]. This heuristic estimates the goal distance by solving a relaxed planning problem in which all the negative effects of actions are ignored. Consider a transportation domain in which trucks are used to move packages between n locations connected in a single chain c_1, \dots, c_n . The goal is to move one package from c_n to c_1 . Figure 1 shows the results of a basic scaling experiment on this domain with $n = 10$ locations, varying the number of trucks T from 1 to 20. All trucks start at c_2 . The results compare basic Monte Carlo Random Walks (MRW) from Arvand-2011 and basic Greedy Best First Search (GBFS) from LAMA-2011. Figure 1 shows how the runtime of GBFS grows quickly with the number of trucks T until it exceeds the memory limit of 64 GB. This is expected since the effective branching factor grows with T . However, the increasing branching factor has only little effect on MRW: the runtime grows only linearly with T .

1.3. Choice of Basic Search Algorithms

All the examples in this paper use state of the art implementations of basic, unenhanced search methods. GBFS as implemented in LAMA-2011 represents systematic search methods, and the MRW implementation of Arvand-2011 represents random walk methods. Both programs use h_{FF} for their evaluation. All other enhancements such as preferred operators in LAMA and Arvand, multi-heuristic search in LAMA, and MHA in Arvand are switched off.

The reasons for selecting this setup are:

1. A focus on theoretical models that can explain the substantially different behavior of random walk and systematic search methods. Using simple search methods allows a close alignment of experiments with theoretical results.

2. Enhancements may benefit both methods in different ways, or be only applicable to one method, so may confuse the picture.
3. A main goal here is to understand the behavior of these two search paradigms in regions where there is a lack of guiding information, such as plateaus. Therefore, in some examples even a blind heuristic is used. While enhancements can certainly have a great influence on search parameters such as branching factor, regress factor, and search depth, the fundamental differences in search behavior will likely persist across such variations.

1.4. Contributions

This paper improves and extends the results reported at the SOCS 2012 conference [14]. The main contributions are:

Regress factor and goal distance for random walks: The key property introduced to analyze random walks is the *regress factor* rf , the ratio of two probabilities: *progressing* towards a goal and *regressing* away from it. Besides rf , the other key variable affecting the average runtime of basic random walks on a graph is the *largest goal distance* D in the whole graph, which appears in the exponent of the expected runtime.

Fair Homogenous graph model: In the *homogenous graph* model, the regress factor of a node depends only on its goal distance and in a *fair* graph a random step changes the goal distance at most by one unit. Theorem 3 shows that the runtime of RW mainly depends on rf . As an example, the state space of Gripper [15] is close to a fair homogenous graph.

Bounds for other graphs: Theorem 4 extends the theory to compute upper bounds on the expected runtime for graphs which are not homogeneous, but for which bounds on the progress and regress chances are known.

Strongly homogenous graph model: In *strongly homogenous graphs*, almost all nodes share the same rf . Theorem 5 explains how rf and D affect the hitting time. A transport example is used for illustration.

Model for Restarting Random Walks: For large values of D , *restarting random walks* (RRW) can offer a substantial performance advantage. At each search step, with probability r a RRW restarts from a fixed initial state s . Theorem 6 gives the ex-




Fig. 1. Average runtime of GBFS and MRW varying the number of trucks (x-axis) in Transport domain. Missing data means memory limit exceeded.

pected runtime of RRW on Homogenous Graphs, relaxing the fairness condition. Furthermore, Theorem 7 proves that the expected runtime of RRW depends only on the goal distance of s , not on D .

Extension to infinitely regressive and non-fair graphs: In *infinitely regressive graphs* and *non-fair* graphs, a random step can arbitrarily increase the goal distance. The main contributions here are Lemma 2 and Theorem 6.

Compared to the conference version, the current paper introduces the extension to infinitely regressive and non-fair graphs. It also contributes more elegant, simpler proofs of Lemma 1 and Theorem 4.

2. Background and Notation

Notation follows standard references such as [16]. Throughout the paper, the notation $P(e)$ denotes the probability of an event e occurring, $G = (V, E)$ is a directed graph, and $u, v \in V$ are vertices.

Definition 1 (Markov Chain). *The discrete-time random process X_0, \dots, X_N defined over a set of states S is Markov(S, \mathbb{P}) iff $P(X_n = j_n | X_{n-1} = j_{n-1}, \dots, X_0 = j_0) = P(X_n = j_n | X_{n-1} = j_{n-1})$. In the matrix $\mathbb{P}(p_{ij})$, $p_{ij} = P(X_n = j_n | X_{n-1} = i_{n-1})$ are the transition probabilities of the chain. In time-homogenous Markov chains as used in this paper, \mathbb{P} does not depend on n .*

Definition 2 (Distance d_G). $d_G(u, v)$ is the length of a shortest path from u to v in G . The dis-

tance $d_G(v)$ of a single vertex v is the length of a longest shortest path from a node in G to v : $d_G(v) = \max_{x \in V} d_G(x, v)$.

Definition 3 (Successors). *The successors of $u \in V$ is the set of all vertices in distance 1 of u : $S_G(u) = \{v | v \in V \wedge d_G(u, v) = 1\}$.*

Definition 4 (Random Walk). *A random walk on G is a Markov chain $\text{Markov}(V, \mathbb{P})$ where $p_{uv} = \frac{1}{|S_G(u)|}$ if $(u, v) \in E$, and $p_{uv} = 0$ if $(u, v) \notin E$.*

The *restarting random walk* model used here is a random walk which *restarts* from a fixed initial state s with probability r at each step, and uniformly randomly chooses among neighbour states with probability $1 - r$.

Definition 5 (Restarting Random Walk). *Let $s \in V$ be the initial state, and $r \in [0, 1]$. A restarting random walk $\text{RRW}(G, s, r)$ is a Markov chain M_G with states V and transition probabilities p_{uv} :*

$$p_{uv} = \begin{cases} \frac{1-r}{|S_G(u)|} & \text{if } (u, v) \in E, v \neq s \\ r + \frac{1-r}{|S_G(u)|} & \text{if } (u, v) \in E, v = s \\ 0 & \text{if } (u, v) \notin E, v \neq s \\ r & \text{if } (u, v) \notin E, v = s \end{cases}$$

A RW is the special case of RRW with $r = 0$.

Definition 6 (Hitting Time). *Let $M = X_0, X_1, \dots, X_N$ be $\text{Markov}(S, \mathbb{P})$, and $u, v \in S$. Let $H_{uv} = \min\{t \geq 0 : X_t = v \wedge X_0 = u\}$. Then the hitting time h_{uv} is the expected number of steps in a random walk on G starting from u which reaches v for the first time: $h_{uv} = E[H_{uv}]$. Therefore, $h_{vv} = 0$.*

Definition 7 (Unit Progress Time). *The unit progress time u_{uv} is the expected number of steps in a random walk after reaching u for the first time until it first gets closer to v . Let $R = \text{RRW}(G, s, r)$. Let $U_{uv} = \min\{t \geq H_{su} : d_G(X_t, v) = d_G(u, v) - 1\}$. Then $u_{uv} = E[U_{uv}]$.*

Definition 8 (Progress, Regress, Infinite Regress and Stalling Chance; Regress Factor). *Let $X : V \rightarrow V$ be a random variable with the following probability mass function:*

$$P(X(u) = v) = \begin{cases} \frac{1}{|S_G(u)|} & \text{if } (u, v) \in E \\ 0 & \text{if } (u, v) \notin E \end{cases} \quad (1)$$

Let X_u be short for $X(u)$. The progress chance $pc(u, v)$, regress chance $rc(u, v)$, infinite regress chance $irc(u, v)$ and stalling chance $sc(u, v)$ of u regarding v , are respectively: the probabilities of getting closer, further away, infinitely further away or staying at the same distance to v after one random step at u .

$$pc(u, v) = P(d_G(X_u, v) = d_G(u, v) - 1)$$

$$rc(u, v) = P(d_G(X_u, v) > d_G(u, v))$$

$$irc(u, v) = P(d_G(X_u, v) = \infty)$$

$$sc(u, v) = P(d_G(X_u, v) = d_G(u, v))$$

The regress factor of u regarding v is $rf(u, v) = \frac{rc(u, v)}{pc(u, v)}$ if $pc(u, v) \neq 0$, and undefined otherwise.

In a Markov Chain, the probability transitions play a key role in determining the hitting time. In all the models considered here, the movement in the chain corresponds to moving between different goal distances. Therefore it is natural to choose progress and regress chances as the main properties.

Theorem 1. [16] *Let M be $\text{Markov}(V, \mathbb{P})$. Then for all $u, v \in V$,*

$$h_{uv} = 1 + \sum_{x \in V} p_{ux} h_{xv}, \quad (2)$$

Theorem 2. *Let $s, u, v \in V$, $R = \text{RRW}(G, s, r)$, $V_d = \{x : x \in V \wedge d_G(x, v) = d\}$, and $P_d(x)$ be the probability of x being the first node in V_d reached by R . Then the hitting time $h_{uv} = \sum_{d=1}^{d_G(u, v)} \sum_{x \in V_d} P_d(x) u_{xv}$.*

Proof. Let the random variable H_{uv} denote the number of steps that R performs since it visits u (for the first time) until it reaches v (for the first time since visiting u), and the random variable X_d denote the first vertex at goal distance d that R reaches after visiting u (Figure 2 shows a schematic representation of these variables). Then

$$H_{uv} = \sum_{d=1}^{d_G(u, v)} \sum_{x \in V_d} 1_{\{X_d\}}(x) U_{xv} \quad (3)$$

where U_{xv} is a random variable measuring the length of the fragment of the walk starting from x




Fig. 2. An illustration of the proof for Theorem 2. Circles represent nodes.

and ending in a smaller goal distance for the first time, and $1_{\{X_d\}}(x)$ is an indicator random variable which returns 1 if $X_d = x$ and 0 if $X_d \neq x$. Since $1_{\{X_d\}}$ and U_{xv} are independent,

$$E[H_{uv}] = \sum_{d=1}^{d_G(u,v)} \sum_{x \in V_d} E[1_{\{X_d\}}(x)]E[U_{xv}]$$

$$h_{uv} = \sum_{d=1}^{d_G(u,v)} \sum_{x \in V_d} P_d(x)u_{xv}$$

2.1. Heuristic Functions, Plateaus, Exit Points and Exit Time

What is the connection between the models introduced here and plateaus in planning? Using the notation of [17], let the heuristic value $h(u)$ of vertex u be the estimated length of a shortest path from u to a goal vertex v . A *plateau* $P \subseteq V$ is a connected subset of states which share the same heuristic value h_P . A state s is an *exit point* of P if $s \in S_G(p)$ for some $p \in P$, and $h(s) < h_P$. The *exit time* of a random walk on a plateau P is the expected number of steps in the random walk until it first reaches an exit point. The problem of finding an exit point in a plateau is equivalent to the problem of finding a goal in the graph consisting of P plus all its exit points, where the exit points are goal states. The expected exit time from the plateau equals the hitting time of this problem. In practice, the search time of planners is often dominated by periods spent in such attempted escapes from plateaus and local minima.




Fig. 3. An illustration of the behaviour of random walks after visiting a node x at the goal distance d .

3. Fair Homogenous Graphs

A fair homogeneous (FH) graph G is the main state space model introduced here. *Homogeneity* means that both progress and regress chances are constant for all nodes at the same goal distance. *Fairness* means that an action can change the goal distance by at most one.

Definition 9 (Homogenous Graph). *For $v \in V$, G is v -homogeneous iff there exist two real functions $pc_G(x, d)$ and $rc_G(x, d)$, mapping $V \times \{0, 1, \dots, d_G(v)\}$ to the range $[0, 1]$, such that for any two vertices $u, x \in V$ with $d_G(u, v) = d_G(x, v)$ the following two conditions hold:*

1. *If $d_G(u, v) \neq 0$, then*
 $pc_G(u, v) = pc_G(x, v) = pc_G(v, d_G(u, v)).$
2. $rc_G(u, v) = rc_G(x, v) = rc_G(v, d_G(u, v)).$

G is homogeneous iff it is v -homogeneous for all $v \in V$. $pc_G(x, d)$ and $rc_G(x, d)$ are called *progress chance* and *regress chance* of G regarding x . The *regress factor* of G regarding x is defined by $rf_G(x, d) = rc_G(x, d)/pc_G(x, d)$.

Definition 10 (Fair Graph). G is fair for $v \in V$ iff for all $u \in V$, for all $x \in S_G(u)$, $|d_G(u, v) - d_G(x, v)| \leq 1$. G is fair if it is fair for all $v \in V$.

Lemma 1. *Let $G = (V, E)$ be FH and $v \in V$. Then for all $x \in V$, h_{xv} depends only on the goal distance $d = d_G(x, v)$, not on the specific choice of x , so $h_{xv} = h_d$.*

Proof. Let $p_d = pc_G(v, d)$, $q_d = rc_G(v, d)$, $c_d = sc_G(v, d)$, $D = d_G(v)$, and $V_d = \{x : x \in V \wedge d_G(x, v) = d\}$. If $d > 0$, then each $x \in V_d$ is connected to at least one node at goal distance $d - 1$.

Thus, $p_d > 0$. The main proof step uses induction from $d + 1$ to d to show that for all $x \in V_d$, $u_{xv} = u_d$. To prove the induction step, assume for all $x' \in V_{d+1}$, $u_{x'v} = u_{d+1}$. The base case for $d = D$ will be shown at the end of the proof since it uses a similar setup as the induction step. After visiting $x \in V_d$ one of the following three cases happens for the random walk (Figure 3):

- with probability p_d it performs a $(d - 1)$ -visit.
- with probability q_d it regresses to the goal distance $d + 1$ and after on average u_{d+1} step it hits $i \in V_d$.
- with probability $1 - p_d - q_d$ it stalls at the same goal distance d hitting $j \in V_d$.

Therefore for $d < D$,

$$u_{xv} = q_d(u_{d+1} + u_{iv}) + (1 - p_d - q_d)u_{jv} + 1$$

The following shows for all $i, j \in V_d$, $u_{xv} = u_d$. Let $\alpha = \arg \max_{k \in V_d} (u_{kv})$ and $\beta = \arg \min_{k \in V_d} (u_{kv})$. Then,

$$\begin{aligned} u_{\alpha v} &= q_d(u_{d+1} + u_{Iv}) + (1 - p_d - q_d)u_{Jv} + 1 \\ &\leq q_d(u_{d+1} + u_{\alpha v}) + (1 - p_d - q_d)u_{\alpha v} + 1 \\ &\leq \frac{q_d}{p_d}u_{d+1} + \frac{1}{p_d} \end{aligned}$$

Furthermore,

$$\begin{aligned} u_{\beta v} &= q_d(u_{d+1} + u_{Iv}) + (1 - p_d - q_d)u_{Jv} + 1 \\ &\geq q_d(u_{d+1} + u_{\beta v}) + (1 - p_d - q_d)u_{\beta v} + 1 \\ &\geq \frac{q_d}{p_d}u_{d+1} + \frac{1}{p_d} \end{aligned}$$

Therefore,

$$\begin{aligned} \frac{q_d}{p_d}u_{d+1} + \frac{1}{p_d} \leq u_{\beta v} \leq u_{xv} \leq u_{\alpha v} \leq \frac{q_d}{p_d}u_{d+1} + \frac{1}{p_d} \\ u_{xv} = \frac{q_d}{p_d}u_{d+1} + \frac{1}{p_d} = u_d \quad (4) \end{aligned}$$

For the base case $d = D$, for all $x \in V_D$

$$\begin{aligned} u_{xv} &= (1 - p_D)u_{Iv} + 1 \\ u_{\alpha v} &\leq \frac{1}{p_D} \\ u_{\beta v} &\geq \frac{1}{p_D} \\ u_{xv} &= \frac{1}{p_D} = u_D \end{aligned}$$

The lemma now follows from Theorem 2:

$$h_{xv} = \sum_{d=1}^{d_G(x,v)} \sum_{k \in V_d} P_d(k)u_{kv} = \sum_{d=1}^{d_G(x,v)} u_d = h_d$$

Theorem 3. Let $G = (V, E)$ be FH, $v \in V$, $p_i = pc_G(v, i)$, $q_i = rc_G(v, i)$, and $d_G(v) = D$. Then for all $x \in V$,

$$h_{xv} = \sum_{d=1}^{d_G(x,v)} \left(\beta_D \prod_{i=d}^{D-1} \lambda_i + \sum_{j=d}^{D-1} \left(\beta_j \prod_{i=d}^{j-1} \lambda_i \right) \right)$$

where for all $1 \leq d \leq D$, $\lambda_d = \frac{q_d}{p_d}$, and $\beta_d = \frac{1}{p_d}$.

Proof. According to Lemma 1 and Theorem 1,

$$\begin{aligned} h_0 &= 0 \\ h_d &= p_d h_{d-1} + q_d h_{d+1} + c_d h_d + 1 \quad (0 < d < D) \\ h_D &= p_D h_{D-1} + (1 - p_D)h_D + 1 \end{aligned}$$

Let $u_d = h_d - h_{d-1}$, then

$$\begin{aligned} u_d &= \lambda_d u_{d+1} + \beta_d \quad (0 < d < D) \\ u_D &= \beta_D \end{aligned}$$

By induction on d , for $d < D$

$$u_d = \beta_D \prod_{i=d}^{D-1} \lambda_i + \sum_{j=d}^{D-1} \left(\beta_j \prod_{i=d}^{j-1} \lambda_i \right) \quad (5)$$

This is trivial for $d = D - 1$. Assume that Equation 5 holds for $d + 1$. Then by Equation 5 for h_{xv} ,

$$\begin{aligned} u_d &= \lambda_d \left(\beta_D \prod_{i=d+1}^{D-1} \lambda_i + \sum_{j=d+1}^{D-1} \left(\beta_j \prod_{i=d+1}^{j-1} \lambda_i \right) \right) + \beta_d \\ &= \beta_D \prod_{i=d}^{D-1} \lambda_i + \lambda_d \sum_{j=d+1}^{D-1} \left(\beta_j \prod_{i=d+1}^{j-1} \lambda_i \right) + \beta_d \\ &= \beta_D \prod_{i=d}^{D-1} \lambda_i + \sum_{j=d+1}^{D-1} \left(\beta_j \prod_{i=d}^{j-1} \lambda_i \right) + \beta_d \prod_{i=d}^{d-1} \lambda_i \\ &= \beta_D \prod_{i=d}^{D-1} \lambda_i + \sum_{j=d}^{D-1} \left(\beta_j \prod_{i=d}^{j-1} \lambda_i \right) \\ h_{xv} &= \sum_{d=1}^{d_G(x,v)} \left(\beta_D \prod_{i=d}^{D-1} \lambda_i + \sum_{j=d}^{D-1} \left(\beta_j \prod_{i=d}^{j-1} \lambda_i \right) \right) \end{aligned}$$

Robot	Gripper	pc	rc	rf	b	d
A	full	$\frac{1}{2}$	$\frac{1}{2}$	1	1	$4 A + 2$
A	empty	$\frac{ A }{ A +1}$	$\frac{1}{ A +1}$	$\frac{1}{ A }$	$ A $	$4 A - 1$
B	full	$\frac{1}{2}$	$\frac{1}{2}$	1	1	$4 A + 1$
B	empty	$\frac{1}{ B +1}$	$\frac{ B }{ B +1}$	$ B $	$ B $	$4 A $

Table 1

Random walks in One-handed Gripper. $|A|$ and $|B|$ denote the number of balls in A and B.

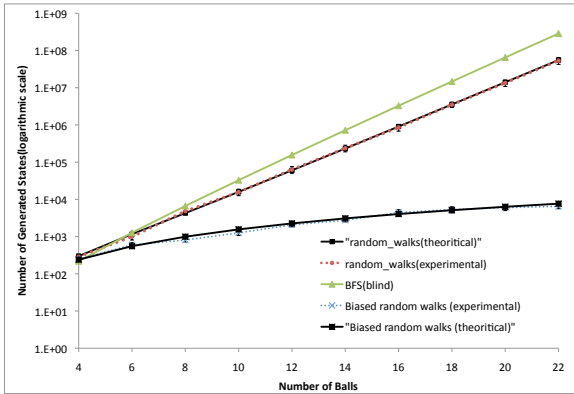


Fig. 4. The average number of generated states varying the number of balls (x-axis) in Gripper domain.

The largest goal distance D and the regress factors $\lambda_i = q_i/p_i$ are the main determining factors for the expected runtime of random walks in homogenous graphs.

3.1. Example domain: One-handed Gripper

Consider a one-handed gripper domain, where a robot must move n balls from room A to B by using the actions of picking up a ball, dropping its single ball, or moving to the other room. The states of the search space fall into four categories shown in Table 1. The search space is fair homogenous: any two states with the same goal distance d have the same distribution of balls in the rooms and also belong to the same category. The graph is fair since no action changes the goal distance by more than one. The expected hitting time is given by Theorem 3.

Figure 4 plots the predictions of Theorem 3 together with the results of a scaling experiment, varying n for both random walks and greedy best first search. To simulate the behaviour of both algorithms in plateaus with a lack of heuristic guidance, a blind heuristic is used which returns 0

for the goal and 1 otherwise. Search stops at a state with a heuristic value lower than that of the initial state. Because of the blind heuristic, the only such state is the goal state. The prediction matches the experimental results extremely well. Random walks outperform greedy best first search. The regress factor rf never exceeds b , and is significantly smaller in states with the robot at A and an empty gripper - almost one quarter of all states.

3.2. Biased Action Selection for Random Walks

Regress factors can be changed by biasing the action selection in the random walk. It seems natural to first select an action type uniformly randomly, then ground the chosen action. In gripper, this means choosing among the balls in the same room in case of the pick up action.

With this biased selection, the search space becomes fair homogenous with $q = p = \frac{1}{2}$. The experimental results and theoretical prediction for such walks are included in Figure 4. The hitting time grows only linearly with n . It is interesting that this natural way of biasing random walks is able to exploit the symmetry inherent in the gripper domain.

4. Extension to Bounds for Other Graphs

While many planning problems cannot be exactly modelled as FH graphs, these models can still be used to obtain upper bounds on the hitting time in any fair graph G which models a plateau. Consider a corresponding FH graph G' with progress and regress chances at each goal distance d respectively set to the minimum and maximum progress and regress chances over all nodes at goal distance d in G . Then the hitting times for G' will be an upper bound for the hitting times in G . In G' , progressing towards the goal is at most as probable as in G .

Theorem 4. *Let $G = (V, E)$ be a fair directed graph, $s, v \in V$, and $D = d_G(v)$. Let $p_{min}(d)$ and $q_{max}(d)$ be the minimum progress and maximum regress chance among all nodes at distance d of v . Let $G' = (V', E')$ be an FH graph, $v', s' \in V'$, $d_{G'}(v') = D$, $pc_{G'}(v', d) = p_{min}(d)$, $rc_{G'}(d) = q_{max}(d)$, and $sc_{G'}(d) = 1 - p_{min}(d) - q_{max}(d)$. Then starting at the same goal distance the hitting time in G' is an upper bound for the hitting time in G , i.e., $h_{sv} \leq h'_{s'v'}$ if $d_G(s, v) = d_{G'}(s', v')$.*

Proof. The first step is to show for all $0 \leq d \leq D$, $sc_{G'}(d) \geq 0$. Let q_x and p_x be the regress and progress chance of node $x \in V$, and $V_d = \{x | x \in V \wedge d_G(x, v) = d\}$, and $j = \arg \max_{x \in V_d} (q_x)$. Then,

$$\begin{aligned} q_{max}(d) &= q_j \leq 1 - p_j \leq 1 - p_{min}(d) \\ q_{max}(d) + p_{min}(d) &\leq 1 \\ sc_{G'}(d) &\geq 0. \end{aligned}$$

Assume for all $x \in V_d$, $u_{xv} \leq u'_d$, where u'_d is the unit progress time at distance d of v' . According to Theorem 2,

$$\begin{aligned} h_{sv} &= \sum_{d=1}^{d_G(s,v)} \sum_{x \in V_d} P_d(x) u_{xv} \\ &\leq \sum_{d=1}^{d_G(s,v)} \sum_{k \in V_d} P_d(x) u'_d \\ &\leq \sum_{d=1}^{d_G(s,v)} u'_d \sum_{k \in V_d} P_d(x) \\ &\leq \sum_{d=1}^{d_{G'}(s',v')} u'_d \\ &\leq h'_d \end{aligned}$$

To prove $u_{xv} \leq u'_d$ by induction, assume for all $x' \in V_{d+1}$, $u_{x'v} \leq u'_{d+1}$ (the induction step; again the base case is shown later). After visiting $x \in V_d$ one of the following three cases happens for the random walk:

- with probability p_x it performs a $(d-1)$ -visit.
- with probability q_x it regresses to the goal distance $d+1$ and, on average, after at least u_{d+1} steps it hits $i \in V_d$.
- with probability $1-p_x-q_x$ it stalls at the same goal distance d hitting $j \in V_d$.

Then for $d < D$,

$$u_{xv} \leq q_x(u_{d+1} + u_{iv}) + (1 - p_x - q_x)u_{jv} + 1.$$

The following shows that for all $i, j \in V_d$, $u_{xv} = u_d$. Let $\alpha = \arg \max_{i \in V_d} (u_{iv})$. Then for $d < D$,

$$\begin{aligned} u_{\alpha v} &\leq q_\alpha(u'_{d+1} + u_{iv}) + (1 - p_\alpha - q_\alpha)u_{jv} + 1 \\ &\leq q_\alpha(u'_{d+1} + u_{\alpha v}) + (1 - p_\alpha - q_\alpha)u_{\alpha v} + 1 \\ &\leq \frac{q_\alpha}{p_\alpha} u'_{d+1} + \frac{1}{p_\alpha} \\ &\leq \frac{q_{max}(d)}{p_{min}(d)} u'_{d+1} + \frac{1}{p_{min}(d)} \end{aligned}$$

Furthermore, according to Equation 4,

$$\frac{q_{max}(d)}{p_{min}(d)} u'_{d+1} + \frac{1}{p_{min}(d)} = u_d \quad (6)$$

Therefore, $u_{xv} \leq u_{\alpha v} \leq u_d$. Analogously, for the base case $d = D$, for all $x \in V_D$

$$u_{\alpha v} \leq \frac{1}{p_\alpha} \leq \frac{1}{p_{min}(d)} \leq u'_d.$$

5. Fair Strongly Homogeneous Graphs

A fair strongly homogenous (FSH) graph G is a FH graph in which pc and rc are constant for all nodes. FSH graphs are simpler to study and suffice to explain the main properties of FH graphs. Therefore, this model is used to discuss key issues such as dependency of the hitting time on largest goal distance D and the regress factors.

Definition 11 (Strongly Homogeneous Graph). *Given $v \in V$, G is strongly v -homogeneous iff there exist two real functions $pc_G(x)$ and $rc_G(x)$ with domain V and range $[0, 1]$ such that for any vertex $u \in V$ the following two conditions hold:*

1. If $u \neq v$ then $pc(u, v) = pc_G(v)$.
2. If $d(u, v) < d_G(v)$ then $rc(u, v) = rc_G(v)$.

G is strongly homogeneous iff it is strongly v -homogeneous for all $v \in V$. The functions $pc_G(x)$ and $rc_G(x)$ are respectively called the progress and the regress chance of G regarding x . The regress factor of G regarding x is defined by $rf_G(x) = rc_G(x)/pc_G(x)$.

Theorem 5. *For $u, v \in V$, let $p = pc_G(v) \neq 0$, $q = rc_G(v)$, $c = 1 - p - q$, $D = d_G(v)$, and $d = d_G(u, v)$. Then the hitting time h_{uv} is:*

$$h_{uv} = \begin{cases} \beta_0 (\lambda^D - \lambda^{D-d}) + \beta_1 d & \text{if } q \neq p \\ \alpha_0 (d - d^2) + \alpha_1 Dd & \text{if } q = p \end{cases} \quad (7)$$

where $\lambda = \frac{q}{p}$, $\beta_0 = \frac{q}{(p-q)^2}$, $\beta_1 = \frac{1}{p-q}$, $\alpha_0 = \frac{1}{2p}$, $\alpha_1 = \frac{1}{p}$.

The proof follows directly from Theorem 3 above. When $q > p$, the main determining factors in the hitting time are the regress factors $\lambda = q/p$ and D ; the hitting time grows exponentially with D and polynomially, with degree D , with λ . As long as λ and D are fixed, changing other structural parameters such as the branching factor b can only increase the hitting time linearly. Note that also for $q > p$, it does not matter how close the start state is to the goal. The hitting time mainly depends on D , the largest goal distance in the graph.

5.1. Analysis of the Transport Example

Theorem 5 helps explain the experimental results in Figure 1. In this example, the plateau consists of all the states encountered before loading the package onto one of the trucks. Once the package is loaded, h_{FF} can guide the search directly towards the goal. Therefore, the exit points of the plateau are the states in which the package is loaded onto a truck. Let $m < n$ be the location of a most advanced truck in the chain. For all non-exit states of the search space, $q \leq p$ holds: there is always at least one action which progresses towards a closest exit point - move a truck from c_m to c_{m+1} . There is at most one action that regresses, in case $m > 1$ and there is only a single truck at c_m which moves to c_{m-1} , thereby reducing m .

According to Theorem 4, setting $q = p$ for all states yields an upper bound on the hitting time, since increasing the regress factor can only increase the hitting time. By Theorem 5, $-\frac{x^2}{2p} + (\frac{2D+1}{2p})x$ is an upper bound for the hitting time. If the number of trucks is multiplied by a factor M , then p will be divided by at most M , therefore the upper bound is also multiplied by at most M . The worst case runtime bound grows only linearly with the number of trucks. In contrast, systematic search methods suffer greatly from increasing the number of vehicles, since this increases the effective branching factor b . The runtime of systematic search methods such as greedy best first search, A* and IDA* typically grows as b^d when the heuristic is ineffective. Regarding the memory usage, since RW in its simplest form only store the *current state* of the walk, increasing the number of trucks does not increase the number of states stored, however, the size of the state grows linearly.

This effect can be observed in all planning problems where increasing the number of objects of a

specific type does not change the regress factor. Examples are the vehicles in transportation domains such as Rovers, Logistics, Transport, and Zenon Travel, or agents which share similar functionality but do not appear in the goal, such as the satellites in the satellite domain. All of these domains contain symmetries similar to the example above, where any one of several vehicles or agents can be chosen to achieve the goal. Other examples are “decoy” objects which can not be used to reach the goal. Actions that affect only the state of such objects do not change the goal distance, so increasing the number of such objects has no effect on rf but can increase b . Techniques such as plan space planning, backward chaining planning, preferred operators, or explicitly detecting and dealing with symmetries can often prune such actions.

Theorem 5 suggests that if $q > p$ and the current state is close to an exit point in the plateau, then systematic search is more effective, since random walks move away from the exit with high probability. This problematic behavior of RW can be fixed to some degree by using restarting random walks.

6. Analysis of Restarting Random Walks

While FH graphs provide bounds for RW on any fair graph, *Infinitely Regressable Homogenous* (IRH) graphs provide bounds for RRW on any strongly homogenous graph. A random step on an IRH graph either gets closer to the goal, stalls at the same goal distance or hits a dead end state, a state with no path to the goal. To keep the analysis simple, it is assumed that all nodes have at least one outgoing edge. Therefore, a RW can never reach a state where no action is available.

Definition 12 (Infinitely Regressable Homogenous Graph). *Given $v \in V$, G is infinitely regressable (IR) v -homogeneous iff for any vertex $u \in V$ there exists at least one vertex x such that $(u, x) \in E$ and there exist three real functions $pc_G(\cdot)$, $sc_G(\cdot)$, and $irc_G(\cdot)$ with domain V and range $[0, 1]$ such that for any vertex $u \in V$ the following three conditions hold:*

1. *If $u \neq v$ then $pc(u, v) = pc_G(v)$.*
2. *$irc(u, v) = irc_G(v)$.*
3. *$sc(u, v) = 1 - irc_G(v) - pc_G(v)$.*




Fig. 5. An illustration of the behaviour of random walks in an IRH graph.

G is IRH iff for any $v \in V$ it is IR v -homogeneous. The functions $pc_G(x)$, $sc_G(x)$ and $irc_G(x)$ are respectively called the progress chance, the stall chance and the infinite regress chance of G regarding x .

Lemma 2. Let $G = (V, E)$ be an IRH graph. Let $RRW(G, s, r)$ be a restarting random walk. Then, for all $v, x, x' \in V$ with $d_G(x, v) = d_G(x', v) = d$ and $d \leq d_G(s, v)$, $h_{xv} = h_{x'v}$.

Proof. Let $p = pc_G(v)$, $c = sc_G(v)$, and $i = irc_G(v)$. Similar to Lemma 1 by induction on the goal distance d , we show that for $d \leq d_G(s, v)$, $u_{xv} = u_d$. Let $V_d = \{x : x \in V \wedge d_G(x, v) = d\}$. Assume for the induction step that for all $x' \in V_{d+1}$, $u_{x'v} = u_{d+1}$. Once more, the proof for the base case follows later. Whenever the random walk transitions to a deadend, it restarts after on average $\frac{1}{r}$ steps (the expected value of a geometric distribution with the success probability r). After each restart a random walk performs on average $U_{d+1} = \sum_{i=d+1}^{d_G(s,v)} u_i$ steps to visit a state with the goal distance d (d -visit). Therefore, after visiting $x \in V_d$ one of the following four cases happens for the random walk (Figure 5):

- with probability r it restarts from s and after on average U_{d+1} steps performs the next d -visit hitting $n \in V_d$.
- with probability $c(1-r)$ it stalls at the same goal distance d hitting $j \in V_d$.
- with probability $i(1-r)$ it transitions to a deadend and after on average $\frac{1}{r} + U_{d+1}$ steps it performs the next d -visit hitting $k \in V_d$.

- with probability $p(1-r)$ it performs a $(d-1)$ -visit.

Therefore, for $d < d_G(s, v)$,

$$u_{xv} = r(U_{d+1} + u_{nv}) + c(1-r)u_{jv} + i(1-r)\left(\frac{1}{r} + U_{d+1} + u_{kv}\right) + (1-r)$$

Note that restarting itself is not counted as a random walk step. The following shows that the identity of nodes n , j and k does not matter. Let $\alpha = \arg \max_{x \in V_d} (u_{xv})$ and $\beta = \arg \min_{x \in V_d} (u_{xv})$. Then,

$$\begin{aligned} u_{\alpha v} &\leq r(U_{d+1} + u_{\alpha v}) + c(1-r)u_{\alpha v} + i(1-r)\left(\frac{1}{r} + U_{d+1} + u_{\alpha v}\right) + (1-r) \\ &\leq \frac{(r + i(1-r))U_{d+1} + (1-r)(1 + i/r)}{(1-r)(1 - i - c)} \end{aligned}$$

Furthermore,

$$\begin{aligned} u_{\beta v} &\geq r(U_{d+1} + u_{\beta v}) + c(1-r)u_{\beta v} + i(1-r)\left(\frac{1}{r} + U_{d+1} + u_{\beta v}\right) + (1-r) \\ &\geq \frac{(r + i(1-r))U_{d+1} + (1-r)(1 + i/r)}{(1-r)(1 - i - c)} \end{aligned}$$

Therefore,

$$\begin{aligned} u_{xv} = u_{\alpha v} = u_{\beta v} = u_d \\ = \frac{(r + i(1-r))U_{d+1} + (1-r)(1 + i/r)}{(1-r)(1 - i - c)} \end{aligned}$$

The base case $d = d_G(s, v)$ has the same four cases, except that after restarting, the random walk immediately performs the d -visit at s :

$$\begin{aligned} u_{xv} &= ru_{sv} + c(1-r)u_{jv} + i(1-r)\left(\frac{1}{r} + u_{kv}\right) + (1-r) \\ u_{xv} = u_{\alpha v} = u_{\beta v} = u_d &= \frac{(1 + i/r)}{(1 - i - c)} \end{aligned}$$

The lemma now follows directly from Theorem 2:

$$h_{xv} = \sum_{d=1}^{d_G(x,v)} \sum_{k \in V_d} P_d(k)u_{kv} = \sum_{d=1}^{d_G(x,v)} u_d = h_d$$

Theorem 6. Let $G = (V, E)$ be an IRH graph, $v \in V$, $p = pc_G(v) > 0$, $c = sc_G(v)$, and $i = irc_G(v)$. Let $R = RRW(G, s, r)$ with $0 < r < 1$. The hitting time $h_{sv} = \Theta(\beta\lambda^{d_s-1})$, where $\beta = \frac{i+r}{rp}$, $\lambda = \frac{i}{p} + \frac{r}{(1-r)p} + 1$ and $d_s = d_G(s, v)$.

Proof. According to Theorem 1 and Lemma 2,

$$\begin{aligned} h_0 &= 0 \\ h_x &= rh_{d_s} + c(1-r)h_x + \\ &\quad i(1-r)\left(\frac{1}{r} + h_{d_s}\right) + (1-r) + ph_{x-1} \end{aligned}$$

Let $u_x = h_x - h_{x-1}$ then

$$\begin{aligned} u_x &= (1-r)(ph_{x-1} - ph_{x-2} + ch_x - ch_{x-1}) \\ &= (1-r)(pu_{x-1} + cu_x) \\ &= \frac{(1-r)p}{1-c+cr} u_{x-1} \end{aligned}$$

Since $c = 1 - p - i$

$$\begin{aligned} u_x &= \frac{(1-r)p}{i(1-r) + p(1-r) + r} u_{x-1} \\ &= \lambda^{-1} u_{x-1} \end{aligned}$$

For $x < d_s$,

$$\begin{aligned} u_x &= \lambda^{d_s-x} u_{d_s} \\ h_x &= \sum_{i=1}^x u_i \\ &= u_{d_s} \sum_{i=1}^x \lambda^{d_s-i} \\ &= \lambda^{d_s-x} \left(\frac{\lambda^x - 1}{\lambda - 1} \right) u_{d_s} \end{aligned}$$

The value u_{d_s} is the progress time from the goal distance d_s . Therefore,

$$\begin{aligned} u_{d_s} &= ru_{d_s} + c(1-r)u_{d_s} + i(1-r)\left(\frac{1}{r} + u_{d_s}\right) + (1-r) \\ &= (r + (1-r)(1-p))u_{d_s} + (i/r + 1)(1-r) \\ &= \frac{i+r}{pr} \\ &= \beta \end{aligned}$$

Therefore,

$$\begin{aligned} h_{d_s} &= u_{d_s} + h_{d_s-1} \\ h_{d_s} &= \beta + \beta\lambda\left(\frac{\lambda^{d_s-1} - 1}{\lambda - 1}\right) \\ h_{d_s} &\in \Theta\left(\beta\lambda^{d_s-1}\right) \end{aligned} \tag{8}$$

The next theorem shows how the results for IRH graphs can be used to derive bounds for any strongly homogeneous graph, even if it is not fair.

Theorem 7. Let $G = (V, E)$ be a strongly homogeneous graph, $v \in V$, $p = pc_G(v) > 0$ and $q = rc_G(v)$. Let $R = RRW(G, s, r)$. The hitting time $h_{sv} \in O(\beta\lambda^{d-1})$, where $\lambda = \left(\frac{q}{p} + \frac{r}{p(1-r)} + 1\right)$, $\beta = \frac{q+r}{pr}$ and $d = d_G(s, v)$.

Proof. For any goal distance x , $h_x \leq \frac{1}{r} + h_d$. This is because the random walk on average restarts from s after $\frac{1}{r}$ steps. The right hand side of this inequality is the hitting time of a random walk stuck in an infinitely large dead end. Therefore, with the pessimistic assumption that each time the random walk regresses from the goal the walk is in a deadend, we can obtain an upper bound for a homogenous graph using the theorem for IRH graphs. It is enough to simply replace i with q in Equation 8.

Therefore, by decreasing r while λ decreases, β increases. Since the upper bound increases polynomially (the degree depends on $d(s, v)$) by λ and only linearly by β , to keep the upper bound low a small value should be chosen for r , especially when $d(s, v)$ is large. The r -value which minimizes the upper bound can be computed from Equation 8.

Comparing the values of λ in the hitting time of RW and RRW, Equations 8 and 7, the base of the exponential term for RRW exceeds the regress factor, the base of the exponential term for RW, by $\frac{r}{p(1-r)} + 1$. For small r , this is close to 1.

The main advantage of RRW over simple random walks is for small $d(s, v)$, since the exponent of the exponential term is reduced from D to $d(s, v) - 1$. Restarting is a bit wasteful when $d(s, v)$ is close to D .

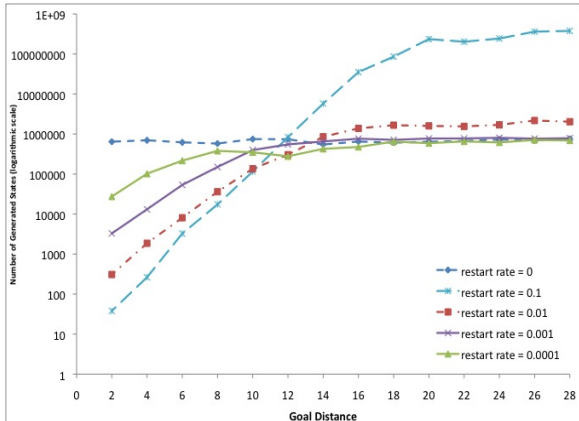


Fig. 6. The Average number of generated states varying the goal distance of the starting state (x-axis) and the restart rate in the Grid domain.

6.1. A Grid Example

Figure 6 shows the results of RRW with restart rate $r \in \{0, 0.1, 0.01, 0.001\}$ in a variant of the Grid domain with an $n \times n$ grid and a robot that needs to first pick up a key at location (n, n) , then unlock a door at $(0, 0)$. The robot can only move left, up or down, except for the top row, where it is also allowed to move right, but not up.

In this domain, all states before the robot picks up the key share the same h_{FF} value. Figure 6 shows the average number of states generated until this subgoal is reached, with the robot starting from different goal distances plotted on the x-axis. Since the regress factors are not uniform in this domain, Theorem 7 does not apply directly. Still, comparing the results of RRW for different $r > 0$ with simple random walks where $r = 0$, the experiment confirms the high-level predictions of Theorem 7: RRW generates slightly more states than simple random walks when the initial goal distance is large, $d \geq 14$, and r is small enough. RRW is much more efficient when d is small; for example it generates three orders of magnitude fewer states for $d = 2$, $r = 0.01$.

7. Related Work

Random walks have been extensively studied in many different scientific fields including physics, finance and computer networking [18,19,20]. Linear algebra approaches to discrete and continuous random walks are well studied [16,21,22,23]. The

current paper mainly uses methods for finding the hitting time of simple chains such as birth–death, and gambler chains [16]. Such solutions can be expressed easily as functions of chain features.

Properties of random walks on finite graphs have been studied extensively [24]. One of the most relevant results is the $O(n^3)$ hitting time of a random walk in an undirected graph with n nodes [25]. However, this result does not explain the strong performance of random walks in planning search spaces which grow exponentially with the number of objects. Despite the rich existing literature on random walks, the application to the analysis of random walk planning seems to be novel.

8. Discussion and Future Work

Important open questions about the current work are how well it models real planning problems such as IPC benchmarks, and real planning algorithms.

Relation to full planning benchmarks: Can they be described within these models in terms of bounds on their regress factor? Can the models be extended to represent the core difficulties involved in solving more planning domains? What is the structure of plateaus within their state spaces, and how do plateaus relate to the overall difficulty of solving those instances? Instances with small state spaces could be completely enumerated and such properties measured. For larger state spaces, can measurements of true goal distances be approximated by heuristic evaluation, by heuristics combined with local search, or by sampling?

Effect of search enhancements: To move from abstract, idealized algorithms towards more realistic planning algorithms, it would be interesting to study the whole spectrum starting with the basic methods studied in this paper up to state of the art planners, switching on improvements one by one and studying their effects under both RW and systematic search scenarios. For example, the RW enhancements MHA and MDA [9] should be studied.

Hybrid methods: Develop theoretical models for methods that combine random walks with using memory and systematic search such as [8,11].

References

- [1] B. Selman, H. J. Levesque, D. Mitchell, A new method for solving hard satisfiability problems, in: Proceedings of the 10th National Conference on Artificial Intelligence, AAAI 1992, San Jose, CA, July 12-16, 1992, 1992, pp. 440–446.
- [2] D. N. Pham, J. Thornton, C. Gretton, A. Sattar, Combining adaptive and dynamic local search for satisfiability, *Journal on Satisfiability, Boolean Modeling and Computation* 4 (2-4) (2008) 149–172.
- [3] S. Gelly, D. Silver, Achieving master level play in 9 x 9 computer Go, in: Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, 2008, pp. 1537–1540.
- [4] H. Finnsson, Y. Björnsson, Simulation-based approach to General Game Playing, in: Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, 2008, pp. 259–264.
- [5] T. Cazenave, Nested Monte-Carlo search, in: Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009, Pasadena, California, USA, July 11-17, 2009, 2009, pp. 456–461.
- [6] C. Browne, E. J. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, S. Colton, A survey of Monte Carlo tree search methods, *IEEE Trans. Comput. Intellig. and AI in Games* 4 (1) (2012) 1–43.
- [7] A. Coles, M. Fox, A. Smith, A new local-search algorithm for forward-chaining planning, in: Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling, ICAPS 2007, Providence, Rhode Island, USA, September 22-26, 2007, 2007, pp. 89–96.
- [8] Q. Lu, Y. Xu, R. Huang, Y. Chen, The Roamer planner random-walk assisted best-first search, in: A. García-Olaya, S. Jiménez, C. Linares López (Eds.), *The 2011 International Planning Competition*, Universidad Carlos III de Madrid, 2011, pp. 73–76.
- [9] H. Nakhost, M. Müller, Monte-Carlo exploration for deterministic planning, in: Proceedings of the 21st International Joint Conference on Artificial Intelligence, IJCAI 2009, Pasadena, California, USA, July 11-17, 2009, 2009, pp. 1766–1771.
- [10] H. Nakhost, J. Hoffmann, M. Müller, Resource-constrained planning: A Monte Carlo random walk approach, in: Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling, ICAPS 2012, Atibaia, São Paulo, Brazil, June 25-19, 2012, 2012, pp. 181–189.
- [11] F. Xie, H. Nakhost, M. Müller, Planning via random walk-driven local search, in: Proceedings of the Twenty-Second International Conference on Automated Planning and Scheduling, ICAPS 2012, Atibaia, São Paulo, Brazil, June 25-19, 2012, 2012, pp. 315–322.
- [12] R. Valenzano, H. Nakhost, M. Müller, J. Schaeffer, N. Sturtevant, ArvandHerd: Parallel planning with a portfolio, in: Proceedings of the 20th European Conference on Artificial Intelligence, ECAI 2012, Montpellier, France, August 27-31, 2012, 2012, pp. 113–116.
- [13] M. Helmert, A planning heuristic based on causal graph analysis, in: Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling, ICAPS 2004, June 3-7, 2004, Whistler, British Columbia, Canada, 2004, pp. 161–170.
- [14] H. Nakhost, M. Müller, A theoretical framework to study random walk planning, in: Proceedings of the Fifth Annual Symposium on Combinatorial Search, SOCS 2012, Niagara Falls, Canada, July 19-21, 2012, 2012.
- [15] D. Long, H. A. Kautz, B. Selman, B. Bonet, H. Geffner, J. Koehler, M. Brenner, J. Hoffmann, F. Rittinger, C. R. Anderson, D. S. Weld, D. E. Smith, M. Fox, The AIPS-98 planning competition, *Artificial Intelligence* 21 (2) (2000) 13–33.
- [16] J. R. Norris, *Markov chains*, Cambridge series in statistical and probabilistic mathematics, Cambridge University Press, 1998.
- [17] H. Hoos, T. Stützle, *Stochastic Local Search: Foundations & Applications*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2004.
- [18] C. Gkantsidis, M. Mihail, A. Saberi, Random walks in peer-to-peer networks: algorithms and evaluation, *Perform. Eval.* 63 (2006) 241–263.
- [19] E. F. Fama, Random walks in stock-market prices, *Financial Analysts Journal* 21 (1965) 55–59.
- [20] H. Qian, S. R. Nassif, S. S. Sapatnekar, Random walks in a supply network, in: Proceedings of the 40th Design Automation Conference, DAC 2003, Anaheim, CA, USA, June 2-6, 2003, 2003, pp. 93–98.
- [21] D. Aldous, J. Fill, *Reversible Markov Chains and Random Walks on Graphs*, University of California, Berkeley, Department of Statistics, 2002.
- [22] G. Yin, Q. Zhang, *Discrete-time Markov chains: two-time-scale methods and applications*, Applications of mathematics, Springer, 2005.
- [23] É. Pardoux, *Markov processes and applications: algorithms, networks, genome and finance*, Wiley series in probability and statistics, Wiley/Dunod, 2009.
- [24] L. Lovász, Random walks on graphs: A survey, *Combinatorics, Paul Erdos is Eighty* 2 (1) (1993) 1–46.
- [25] G. Brightwell, P. Winkler, Maximum hitting time for random walks on graphs, *Random Struct. Algorithms* 1 (1990) 263–276.