

Stochastic Search With Locally Clustered Targets: Learning from T Cells

Rüdiger Reischuk and Johannes Textor

Institut für Theoretische Informatik
Universität zu Lübeck, 23538 Lübeck, Germany
{reischuk,textor}@tcs.uni-luebeck.de

Abstract. Searching a space with locally clustered targets (think picking apples from trees) leads to an optimization problem: When should the searcher leave the current region, and invest the time to travel to another one? We consider here a model of such a search process: infection screening by T cells in the immune system. Taking an AIS perspective, we ask whether this model could provide insight for similar problems in computing, for example Las Vegas algorithms with expensive restarts or agent-based intrusion detection systems. The model is simple, but presents a rich phenomenology; we analytically derive the optimal behavior of a single searcher, revealing the existence of two characteristic regimes in the search parameter space. Moreover, we determine the impact of perturbations and imprecise knowledge of the search space parameters, as well as the speedup gained by searching in parallel. The results provide potential new directions for developing tools to tune stochastic search algorithms.

1 Introduction

Natural resources such as fruit, drinking water, minerals, or prey are most often unevenly distributed in the environment. Moreover, these resources are not infinite, and may be depleted by consumption. Animals thus need to adjust their foraging behaviour accordingly. For example, many predators migrate to different hunting grounds from time to time to ensure a continuous supply with prey. *Foraging theory* [1] is a mathematical treatment of animal foraging behaviour based on the hypothesis that animals evolve to maximize their energy intake, and thus find ways to use their environment optimally. This gives rise to optimization problems such as the *giving up time*: when should a forager give up its current hunting ground, and invest the energy necessary to find a new one?

In this paper, we introduce a “foraging-type” model that we developed to understand the stochastic search of T cells for antigen. Taking an AIS perspective, we ask: what could we learn from this model that could be useful for similar optimization processes in computing and operations research? Scenarios that lead to optimization problems of this type include the following:

- *Security*: Immune-inspired distributed intrusion detection systems (e.g. [2], [3]) mimicking the function of T cells: They consist of large numbers of

- agents, each specialized to detecting a certain type of intrusion, which continuously migrate between different hosts or switches in the network.
- *Economy*: Quality control procedures in companies that are organized into branches. The employees in charge of quality control divide their time between inspecting branches and travelling between branches.
 - *Algorithms*: Optimal restart of Las Vegas algorithms with expensive restarts. A Las Vegas algorithm is an algorithm with deterministic output whose run-time is a random variable. Such algorithms can be accelerated by restarting them when an unfavorable region of the run time distribution is reached; a well-known example is Schönig’s probabilistic SAT solver [4]. An optimal strategy for restarting such algorithms was given by Luby et al. [5]. However, this strategy does not take into account the time needed for restarting the algorithm, which may be substantially larger than the cost of a single search step [4]. Taking this into account leads to a “foraging-type” optimization problem.

1.1 Model Definition and Contributions

The general case of our model can be stated as follows. Consider an indexed set of compartments, which we call *bags*. For each bag i , we have a discrete probability density function $P_i(t)$, giving the probability of hitting the target at time t in bag i , where $t = 0$ is the time that the bag is entered. Furthermore, we have a function $T(i, j)$ denoting the time needed to travel from bag i to bag j . We will assume that the searcher has no knowledge about the *local* target densities (it does not know the $P_i(t)$ for the bags i it visits), and thus simply chooses the next bag uniformly at random. However, the searcher may have some knowledge about the *global* target density (e.g., a suitable average of all the $P_i(t)$).

A searcher like this learns nothing as it progresses, because the search is already finished upon encountering the target. As proved by Luby et al. [5], the optimal strategy in such a case is to restart the search in fixed time intervals. A *strategy* can thus be defined simply as a constant *residence time* R to be spent in each bag before travelling to a different one. Denote by H the *first hitting time*, i.e., the random variable giving the time that the searcher finds the first target. Our goal is to set R such as to minimize the expected hitting time $E[H]$. In more complex cases, one might consider R as a function whose input reflects the searcher’s knowledge about its environment.

In the present paper we focus on a special case of our model where the transit time $T(i, j)$ is constant, and sampling in the bags is by drawing with replacement. This special case is a reasonable model of T cell infection screening in the immune system, which is justified in detail in Section 1.2. Readers who are not interested in the immunological background can skip the rest of this introduction and move directly to the technical part, which starts in Section 2. Our technical contributions are the following:

- We obtain both exact and asymptotic results for a single searcher’s optimal behaviour (Section 3). This analysis clearly identifies two distinct parameter regimes with quite different asymptotics.

- We then apply these results to characterize the impact of parameter perturbations on the search performance (Section 3.5). This gives insight into the robustness of an optimal searcher against such perturbations, as well as bounds for the search performance when only imprecise estimates of the search space parameters are available.
- Finally, we analyze the performance of a parallel search by multiple independent searchers (Section 4).

1.2 Background: Modelling T Cell Immune Surveillance

The T cells of the immune system [6] screen for antigen (e.g. viruses, bacteria) presented on specialized cells in lymph nodes. A mouse, for instance, has 30-35 lymph nodes distributed strategically across its body. Detection of antigen by T cells is an important step for the initiation of an immune response. T cells are highly specialized: It was estimated that for a given antigen, a mouse has only around 100-200 T cells capable of detecting that antigen [7]. Because many infections are localized (e.g. in the respiratory tract or the intestine), T cells continuously circulate around the body in search for antigen, and migrate to a different lymph node approximately once per day [8]. The cells are carried to the lymph nodes via the blood stream, and hence essentially circulate at random. Mapping this to the previously defined notation, the bags are the lymph nodes, R corresponds to the time spent in the lymph node, and $T(i, j)$ becomes a constant T equal to the time spent travelling between lymph nodes. Since T cells spend around 2/3 of their lifetime in lymph nodes [9], $R \approx 18h$ and $T \approx 6h$.

The most interesting question is whether drawing with replacement is an accurate model for antigen sampling within lymph nodes. By means of two-photon microscopy, it has become possible to observe T cells in lymph nodes in the living, intact animal. These experiments revealed that the search of T cells for antigen is essentially a 3D random walk through the lymph node tissue [10]. Thus, we can indeed justify the drawing with replacement model as it is a quite reasonable approximation of a 3D random walk. In the rest of this section we give some more detail about this correspondence (note however, that this is a mere exposition of well-known facts from random walk theory).

Consider a random walk in the lattice \mathbb{Z}^3 , and let p_n denote the probability that the vertex reached by the random walk in its n -th step has been visited before. S_n denotes the expected number of different vertices covered within the first n steps, and r_n the probability that the random walk returns to its point of origin at least once within the first n steps. By reversing the random walk in time, it is clear that $p_n = r_n$, and it is known that

$$r_\infty := \lim_{n \rightarrow \infty} r_n = \sup_{n \in \mathbb{N}} r_n = 1 - \frac{1}{u(3)} = 0.3405373\dots, \quad (1)$$

which is Polya's random walk constant [11]. Thus, we obtain the following asymptotic bound on S_n :

$$E[S_n] \approx (1 - r_\infty) n \quad (2)$$

Hence, every time a node is visited, we have a chance of at least 66% that it is a node we have not visited before. Now assume that targets are distributed in this lattice such that every vertex is a target with probability c . We are interested again in the first hitting time, i.e. the number of steps it takes a random walk to hit a target for the first time, denoted by $H = H(c)$. Its expectation can be evaluated as follows:

$$\mathbb{E}[H] = \sum_{j=1}^{\infty} j \cdot \Pr[H = j] = \sum_{j=1}^{\infty} j \cdot (\Pr[H > j-1] - \Pr[H > j]) \quad (3)$$

$$= \sum_{j=0}^{\infty} \Pr[H > j] = \sum_{j=0}^{\infty} \mathbb{E}[(1-c)^{S_j}] \quad (4)$$

$$(5)$$

For $c \ll 1$, i.e. when $(1-c)^n$ decreases very slowly, the last expression can be approximated by

$$\mathbb{E}[H] \approx \sum_{j=0}^{\infty} (1-c)^{\mathbb{E}[S_j]} \approx \sum_{j=0}^{\infty} (1-c)^{(1-r_{\infty})j} \quad (6)$$

$$= \frac{1}{1 - (1-c)^{(1-r_{\infty})}} = \frac{1}{1 - (1-c)^{0.659462670\dots}} \quad (7)$$

This approximation, called the *Rosenstock approximation* [12], is known to give quite good estimates for $c < 0.05$. Using the binomial series expansion, and setting $t_{\infty} = 1 - r_{\infty}$, we obtain for the (per step) success probability of the random walk:

$$\frac{1}{\mathbb{E}[H]} \approx 1 - (1-c)^{t_{\infty}} = 1 - \sum_{k=0}^{t_{\infty}} \binom{t_{\infty}}{k} (-c)^k = c t_{\infty} - O(c^2) \quad (8)$$

Comparing this random walk search to a systematic search of the lattice vertex by vertex (which is equivalent to drawing with replacement), we see that both methods give rise to a geometric hitting time distribution. For small c , the systematic search outperforms the random walk in terms of expected hitting time by a factor of at most 1.5. Note that at least three spatial dimensions are needed to make the random walk search competitive with the systematic search – the random walk search strategy would be far less effective in a one- or two-dimensional environment.

2 Formal Statement of the Special Case

Recall the definitions of *bags* with according hitting time distributions $P_i(t)$ and transit times $T(i, j)$ from Section 1.1. We assume that the local search in a bag is by sampling with replacement, i.e., for all bags i we have $P_i(t) = c_i$. Hence, the hitting time within bag i is geometrically distributed with parameter c_i . We

distinguish between two types of bags, called *good* bags and *bad* bags¹. We fix some constant c such that $c_i = c$ for all good bags c_i . For the bad bags c_j , we set $c_j = 0$. Furthermore, we assume the travel time $T(i, j) = T$ to be constant. Let n_{good} denote the number of good bags, and n be the number of all bags. Then $\nu := n_{\text{good}}/n$ gives the fraction of good bags². Hence, when sampling the bag i , the success probability in the current bag (which was drawn uniformly at random from $\{1, \dots, n\}$) is c with probability ν and 0 with probability $1 - \nu$. Because the next bag to travel to is chosen at random, our special case is now fully characterized by the parameters c, T , and ν , and we will only consider the nontrivial cases with $0 < \nu < 1$.

As mentioned in Section 1.1, it follows from Luby et al. [5] that the optimal strategy is to transit to a new bag after a fixed residence time R_{opt} . Even though we have yet to determine R_{opt} , this means that we can restrict our attention to searchers of the following type: The searcher picks a bag i uniformly at random, spends time R in the bag drawing one sample per time step, and then spends time T in transit to the next bag. This is iterated until a target is found. Despite its simplicity, our special case gives rise to surprisingly rich asymptotic behaviour, as we will see in the upcoming analysis.

3 Optimizing a Single Searcher

Since the parameters T, ν and c characterize the search problem and are thus beyond our control, our goal is to tune the parameter R in order to maximize search performance. This leads to an optimization problem: If a searcher spends too much time in a bag i , it risks that i is a bad bag, in which case the time would better be spent searching somewhere else. On the other hand, leaving the bag i carries the risk that i could have been a good bag and the target might soon have been found. If the searcher could detect which type of bag it is currently in, then the optimal strategy would obviously be to transit immediately to new bags until a good one is reached. However, because this is not possible, R must be set to a value that balances between the two risks. We derive the following results analytically:

- Expectation $E[H]$ of the hitting time H (Proposition 1);
- Asymptotics of $E[H]$ for large and small R (Propositions 2,3);
- The optimal residence time R_{opt} (Proposition 4); and
- Asymptotic expressions for R_{opt} for *locally dominated* and *globally dominated* parameters (Propositions 5 and 6).

¹ Note that there must at least be two different types of bags, otherwise there would of course be no benefit in travelling between bags.

² In the immune surveillance model, ν corresponds to the fraction of lymph nodes that are near the area of infection (draining lymph nodes). For instance, in a literature experiment with herpes simplex virus [13], there were 5 draining lymph nodes on average, thus $\nu = 5/35$.

3.1 The Expected Hitting Time

Proposition 1. *Let H denote the first hitting time of a search process according to Section 2 with parameters R, T, c , and ν . Then H has expectation*

$$\mathbb{E}[H] = \frac{1-c}{c} - \frac{1-\rho}{\rho} R + \left(\frac{1-q}{q}\right) (R+T) + 1 \quad (9)$$

where $\rho = 1 - (1-c)^R$ and $q = \nu\rho$.

Proof. Let us call a sequence of searching a bag (which takes time R) and transiting to the next bag (which takes time T) a *phase*. Let U be a random variable denoting the number of unsuccessful phases before the searcher finds a target in a good bag in phase $U+1$, and let S be the number of samples drawn in phase $U+1$ before the target is found. Then the hitting time is given by

$$H = (T+R)U + S + 1. \quad (10)$$

Since U and S are stochastically independent, it holds that

$$\mathbb{E}[H] = (R+T) \mathbb{E}[U] + \mathbb{E}[S] + 1. \quad (11)$$

U is geometrically distributed with parameter q , hence $\mathbb{E}[U] = (1-q)/q$. S on the other hand has a geometric distribution that is “truncated” to the finite support $\{0, \dots, R-1\}$. With some algebra, it can be verified that

$$\mathbb{E}[S] = \frac{1}{\rho} \sum_{k=0}^{R-1} k (1-c)^k c = \frac{1-c}{c} - \frac{1-\rho}{\rho} R. \quad (12)$$

Putting S and U together, we obtain the result. \square

The above formula is indicative of the fact that the search is a combination of two sampling processes with replacement: The global search for a good bag, and the local search for a target in a good bag.

3.2 Asymptotics of the Expected Hitting Time

To understand the dependencies of the expected hitting time, we first analyze its asymptotics for large and small R . For $R \geq M c^{-1}$ with $M \gg 1$ the term $(1-c)^R \leq e^{-M}$ becomes very small, and thus ρ close to 1. This results in the following asymptotics:

Proposition 2. *Let H, R, T, ν, c be defined as above, and fix a large constant $M \gg 1$ such that $R \geq M/c$. Then*

$$\mathbb{E}[H] \approx \frac{1-c}{c} + \frac{1-\nu}{\nu} (R+T) \in \Theta(R). \quad (13)$$

Hence, spending significantly more time searching a bag than the expected hitting time c^{-1} for a good bag increases the overall hitting time linearly. On the other hand, if R becomes too small, we get the following:

Proposition 3. *Let H, R, T, ν, c be defined as above, and fix some small nonzero $\epsilon \ll 1$ such that $R \leq \epsilon c^{-1}$. Then*

$$E[H] \approx \frac{1}{\nu c} + \frac{T}{\nu c} R^{-1} - T \in \Theta(1/R). \quad (14)$$

Proof. If $R \leq \epsilon c^{-1}$, we can use the approximation $(1 - c)^R = 1 - R c + O((R c)^2)$. This implies $\rho \approx c R$ and $q = \nu c R$, which upon insertion into Equation 9 gives the result. \square

The $\Theta(1/R)$ asymptotics for small R (i.e., halving an already small R almost doubles the number of phases until a global hit occurs) can be intuitively explained by noting that most of the time is spent in transit between bags since the success probability within a bag $\approx R/(\nu c)$ is very low (Figure 1).




Fig. 1. The expected hitting time $E[H]$ as per Proposition 9 and its asymptotics as per Propositions 2 and 3 as a function of the residence time R for the parameters $\nu = 0.1, T = 100, c = 0.001$.

3.3 The Optimal Residence Time

For given ν, T and c , what is the optimal choice for R , i.e., the one that minimizes $E[H]$? Let us denote this value by R_{opt} . It is given by the following proposition:

Proposition 4. *Let H, R, T, ν, c be defined as above, and consider $E[H]$ as a function of R where T, ν, c are constant. Then the $E[H]$ is minimized by*

$$R_{opt} = W_{-1} \left(-\frac{(1-c)^{\frac{T}{1-\nu}}}{e} \right) \frac{1}{\ln(1-c)} - \frac{T}{1-\nu} + \frac{1}{\ln(1-c)} \quad (15)$$

where W_{-1} is the non-principal branch of the Lambert W function [14].

Proof. We have to solve $\frac{d}{dR} E[H] = 0$, which is equivalent to

$$0 = \frac{d}{dR} \left[\frac{1-\rho}{\rho} R + \frac{1-q}{q} R + \frac{1}{q} T \right] = \frac{d}{dR} \left[\left(\frac{1}{\nu} - 1 \right) \frac{R + T/(1-\nu)}{1 - (1-c)^R} \right]. \quad (16)$$

Now, differentiating we get

$$0 = 1 - (1-c)^R + (R + T/(1-\nu)) (1-c)^R \ln(1-c) \quad (17)$$

$$\Leftrightarrow \left(R + \frac{T}{1-\nu} - \frac{1}{\ln(1-c)} \right) (1-c)^R = -\frac{1}{\ln(1-c)}. \quad (18)$$

This is a transcendental equation and thus cannot be solved for R using only standard algebra. A tool for solving equations of this type, which arise in many applications [14], is the *Lambert W function* defined by

$$x e^x = y \iff x = W(y). \quad (19)$$

Using this function, we can express the solutions \hat{x} of the equation $(\hat{x} + \beta) \alpha^{\hat{x}} = \gamma$ in closed form with $\hat{y} = \alpha^\beta \gamma \ln \alpha$, because $\hat{x} = \frac{W(\hat{y})}{\ln \alpha} - \beta$. Inserting $\alpha = (1-c)$, $\beta = \frac{T}{1-\nu} - \frac{1}{\ln(1-c)}$ and $\gamma = -\frac{1}{\ln(1-c)}$ in our case gives $\hat{y} = -(1-c)^{\frac{T}{1-\nu}}/e$. Because $-1/e < \hat{y} < 0$, the two branches W_0 and W_{-1} of the Lambert W function both solve the equation. The non-principal branch W_{-1} is the meaningful one in our case because it maps to the interval $(-\infty, -1)$, while W_0 maps to $[-1, 0]$. Inserting α, β, γ and \hat{y} yields the claimed expression. \square

3.4 Asymptotics of the Optimal Residence Time

The exact solution given by Proposition 4 for the optimal residence time is rather complex and yields little insight into the dependencies between R_{opt} and the parameters T, ν , and c . Thus, we now turn our attention to two important regions in the parameter space for which more illustrative asymptotic forms of R_{opt} can be derived. For simplicity, we assume that c is moderately small (e.g. $c < 0.1$); a similar analysis is possible without this assumption, but the asymptotic formulae become more complicated.

We will show the existence of two quite different parameter regimes. The switching point between them is given by a rather unexpected trade-off:

Definition 1. Let H, T, ν, c be defined as above, and let c be moderately small such that $\ln(1-c) \approx -c$. For

$$\frac{1}{c} \ll \frac{T}{1-\nu}, \quad (20)$$

we call H transit dominated. Otherwise, if

$$\frac{1}{c} \gg \frac{T}{1-\nu}, \quad (21)$$

then we call H locally dominated.

Note that the parameter ν plays hardly any role in defining these two parameter regimes as in the interesting cases, ν is typically rather small (otherwise the search problem would not be very different from a simple local search in one bag). Thus, surprisingly, *the transit time T is more important than the difficulty of the global search problem*, which can be measured by $1/\nu$ rather than $1/(1-\nu)$.

The upcoming two propositions yield quite interesting differences between transit dominated and locally dominated settings.

Proposition 5. *Let H be transit dominated by T, ν, c . Then*

$$R_{opt} \approx \frac{\ln T - \ln(1-\nu) + \ln c}{c}. \quad (22)$$

Proof. We use the following power series expansion for $W_{-1}(y)$, which converges quickly for $1/e \ll y < 0$ [14]:

$$W_{-1}(y) = \lambda_1 - \lambda_2 + \sum_{k=0}^{\infty} \sum_{m=1}^{\infty} c_{km} \frac{\lambda_2^m}{\lambda_1^{m+k}} \quad (23)$$

$$= \lambda_1 - \lambda_2 + \sum_{k=0}^{t-1} \sum_{m=1}^{t-k} c_{km} \frac{\lambda_2^m}{\lambda_1^{m+k}} + O\left(\left(\frac{\lambda_2}{\lambda_1}\right)^{t+1}\right), \quad (24)$$

where $\lambda_1 := \ln(-y)$, and $\lambda_2 := \ln(-\lambda_1) = \ln(-\ln(-y))$. The c_{km} are constants that are not important for our analysis, since we asymptotically approximate W_{-1} for $y \rightarrow 0$ by truncating the sum terms of the power series ($t=0$). For our \hat{y} defined in the proof of Proposition 4 this results in $\lambda_1 = \frac{T}{1-\nu} \ln(1-c) - 1$ and

$$W_{-1}\left(-\frac{(1-c)^{\frac{T}{1-\nu}}}{e}\right) = \lambda_1 - \lambda_2 + O\left(\frac{\lambda_2}{\lambda_1}\right) \quad (25)$$

$$= \frac{T \ln(1-c)}{1-\nu} - 1 - \ln\left(1 - \frac{T \ln(1-c)}{1-\nu}\right) + O\left(\frac{\lambda_2}{\lambda_1}\right). \quad (26)$$

Inserting this asymptotic expansion into the closed form for R_{opt} given by Proposition 4, some terms cancel out and we arrive at

$$R_{opt} = -\ln\left(1 - \frac{T \ln(1-c)}{1-\nu}\right) \frac{1}{\ln(1-c)} \quad (27)$$

$$+ O\left(\frac{\lambda_2}{\lambda_1 \ln(1-c)}\right). \quad (28)$$

In the region where $\frac{T}{1-\nu} \gg -\ln(1-c) \approx c$, the argument $1 - \frac{T \ln(1-c)}{1-\nu}$ of the first logarithm is much larger than 1 and can be replaced by $\frac{T}{1-\nu} (-\ln(1-c))$. This gives the approximation

$$R_{opt} \approx \frac{\ln T - \ln(1-\nu) + \ln(-\ln(1-c))}{-\ln(1-c)}. \quad (29)$$

which is valid for any c in a transit dominated setting. Substituting $-\ln(1-c)$ for c yields the claimed expression. \square

To understand why this approximation eventually breaks down for $c \rightarrow 0$, we look more closely at the O -term of Equation (28) for R_{opt} :

$$R_{\text{opt}} = -\ln\left(1 - \frac{T \ln(1-c)}{1-\nu}\right) \frac{1}{\ln(1-c)} + O\left(\frac{\ln\left(1 - \frac{T \ln(1-c)}{1-\nu}\right)}{\frac{T \ln(1-c)^2}{1-\nu} - \ln(1-c)}\right) \quad (30)$$

Applying De l'Hôpital's Rule it can be shown that R_{opt} *without* the O -term approaches a constant value for $c \rightarrow 0$, whereas the O -term starts to dominate. This limit takes us to the locally dominated regime, which we examine next.

Proposition 6. *Let H be locally dominated by T, ν, c . Then we have*

$$R_{\text{opt}} \approx \sqrt{\frac{2T}{(1-\nu)c}}. \quad (31)$$

Proof. In the transit dominated regime, the argument of W_{-1} in the closed form of R_{opt} (Proposition 4) is close to $-1/e$, the branch point of the W function. Near this branch point, the power series used in the proof of the previous result is no longer useful (Figure 2, left). From the results of Corless et al. [14] one can derive an alternative power series expansion for W_{-1} near the branch point:

$$W_{-1}(y) = \sum_{t=0}^{\infty} (c_t)^t = -1 + \sigma - \frac{1}{3}\sigma^2 + \frac{11}{72}\sigma^3 + \dots \quad (32)$$

In this expression, $\sigma = -\sqrt{2ey+2}$, and thus $|\sigma| \leq 1$. Again the c_t are constants that are irrelevant for our purpose, since we truncate the series after $t = 1$ to obtain an asymptotic approximation. Inserting again the argument for y yields

$$W_{-1}\left(-\frac{(1-c)^{\frac{T}{1-\nu}}}{e}\right) = -1 - \sqrt{2 - 2(1-c)^{\frac{T}{1-\nu}}} + O\left(1 - (1-c)^{\frac{T}{1-\nu}}\right) \quad (33)$$

from which we get the following expression for R_{opt} :

$$R_{\text{opt}} = \frac{-\sqrt{2}}{\ln(1-c)} \sqrt{1 - (1-c)^{\frac{T}{1-\nu}} + O\left(\left(1 - (1-c)^{\frac{T}{1-\nu}}\right)^2\right)}. \quad (34)$$

By the definition of locally dominated parameters, we have $Tc/(1-\nu) \ll 1$. Thus we can substitute $(1-c)^{T/(1-\nu)}$ by $1 - Tc/(1-\nu)$. This yields the claimed expression.

3.5 Implications for Robustness and Parameter Estimation

The asymptotic results derived in the previous section yield important insight for situations where we have no exact knowledge about the parameters of the search




Fig. 2. Left: Illustration of the two different approximations (Equations 23 and 32) used for the Lambert W function in the proofs of Propositions 5 and 6. Right: Transition of the optimal residence time between the two regions described by Proposition 6 (densely dashed) and Proposition 5 (dashed) for $T = 1000, \nu = 0.1$ and varying c . The square marks the point where $\frac{T}{1-\nu} = \frac{1}{c}$.

problem. For example, consider the following two questions: (1) A searcher’s residence time has been optimally calibrated, and now one of the search parameters is perturbed. How much would the perturbation affect the searcher’s performance? (2) We determine the search space parameters by statistical estimation. How precise would our estimate need to be to get reasonable performance?

Assume that we set our R to within a factor κ of R_{opt} , i.e., $R_{\text{opt}}/\kappa < R < \kappa R_{\text{opt}}$. Then it follows from the results in Sections 3.1 and 3.2 that $E[H]$ is also within a factor κ of its optimal value. Combining this with the results from the previous section, we see that the situation is very different for the two parameter regimes:

In the *locally dominated* regime (Proposition 6), we have *square root* asymptotics for $1/c, 1 - \nu$, and R , implying that $E[H]$ would be within factor $\sqrt{\kappa}$ of its optimal value if one of these parameters is perturbed by factor κ . Hence, the perturbation has sublinear impact. In the *transit dominated* regime (Proposition 5), the effect of perturbing T and $1 - \nu$ would even be merely *logarithmic*; however, the effect of perturbing c in this regime would be *linear*. Note that in either case, perturbing ν instead of $1 - \nu$ by a small factor κ has virtually no effect if ν is already small.

4 Parallel Search

In this section, we ask how the expected hitting time is reduced by employing several searchers in parallel. We will limit our discussion to the case that the searchers are not synchronized. Note however that being synchronized or not only makes a substantial difference if both ν and c are close to 1, in which case the search problem is anyway not very difficult. We assume that the searchers cannot communicate with each other, and that there is no global control.

Let m denote the number of searchers and H_m the hitting time of such a parallel search. For values of m that are significantly smaller than the expected

hitting time $E[H]$ of a single agent, the expectation of the m -parallel search can be approximated by

$$E[H_m] \approx \frac{E[H]}{m}, \quad (35)$$

since in this case the hitting probability of a single step grows approximately by a factor m . This approximation will become invalid for large m , because the bags become saturated with searchers and thus additional searchers will no longer yield substantial speedup. However, in this situation it is still possible to use the following approximation instead:

$$E[H_m] = \frac{1}{1 - (1 - R\nu c / (R + T))^m} \quad (36)$$

This approximation is obtained by noting that for a randomly chosen time step, every searcher has an overall chance of $R\nu/(R + T)$ to be in a good bag, and within a good bag the chance of finding a target is c . Assuming that the fraction of searchers in good bags at every timestep is indeed equal to $\gamma = R\nu/(R + T)$ (rather than a random variable with expectation γ), we can approximate the parallel search by random sampling with replacement with a success probability of $1 - (1 - \gamma c)^m$. Note that for $m = 1$, the above equation is equal (up to the constant T) to the equation in Proposition 3 describing the asymptotics of $E[H]$ for small R . Some experimental results are displayed in Figure 3. Notably, while for optimally tuned searchers the speedup is indeed well described by the above equations that predict a linear speedup with saturation, for *non-optimal* residence times it is possible to obtain a *superlinear* speedup by increasing the number of searchers (e.g. in Figure 3, around $m = 100$ for $R = 1$ and around $m = 20$ for $R = 100$).




Fig. 3. Speedup of the expected hitting time by m independent parallel searchers. $E[H]$ is plotted as a function of m for $T = 10, \nu = 0.1$, and $c = 0.12591$, which gives $R_{\text{opt}} = 10$. For the optimally tuned population, the approximate expected hitting times as predicted by Equations 35 and 36 are displayed, where Equation 35 describes a power law with slope -1. Per data point we performed 1000 simulations, so that all standard errors of the mean were less than 2%.

5 Conclusions and Future Work

We have stated a “foraging-type” model of stochastic search with locally clustered targets. We performed an in-depth theoretical analysis of the optimal behaviour of a single searcher in this model. The special case corresponds to the behaviour of T cells in the immune system, and is more generally valid for all search processes of this type where the local search is reasonably approximable by drawing with replacement (e.g., the local search is a ≥ 3 D random walk). As mentioned in the beginning, optimization questions of this type arise in many fields, and we are confident that the model can be applied or extended to many such problems as it is not very immune system-specific. Our work raised many questions to be pursued in the future:

Generalization to other hitting time distributions. The most obvious question is how we can generalize the analytical results obtained for our special case to other scenarios. For instance, concerning our application of modelling infection screening by T cells, there is a hypothesis [15] that the cells need to hit several targets in order to be activated rather than just one. This would give rise to the following generalization of our special case: If c is the probability for a single hit in a good bag per time unit, and k is the number of hits, then the expected hitting time within a good bag would follow a *negative binomial distribution* with parameters c and k instead of a geometric one (which is equal to a negative binomial distribution with $k = 1$). Other distributions of interest include the Weibull distribution which can be used to model the time to failure of systems with constant failure rate; this could be interesting for the quality control scenario.

Parameter estimation and adaptive search. Often we will not know the search space parameters precisely. Thus, the search cannot guarantee to start with an optimal behaviour for the given system. How does a good or optimal strategy look like to estimate the parameters and to adjust the search to the optimal values? Similarly, the parameters may change over time. For instance, an intrusion detection system might face a sequence of intrusions where each intrusion is governed by a certain parameter setting. This setting may change slowly over the time. Knowing the parameters at the beginning or after some period of estimation the searchers could use a strategy optimized for this setting. But later after the change of some parameters this might not longer be the case. How could one adjust in this case?

Application to combinatorial optimization problems. In the introduction, we mentioned the connection of our model to Las Vegas algorithms with expensive restarts. Our model is potentially applicable to combinatorial optimization problems having the property that solutions form clusters in the search space, which e.g. is the case for suitably encoded versions of the traveling salesman problem (TSP). Large instances of the TSP would also have the property that

generating a starting point (a tour with a reasonably low weight) takes substantially longer than a local search step in the problem space (e.g. by locally modifying some edges). It remains to be seen whether taking the restart time into account explicitly could lead to substantially faster stochastic algorithms for such problems. In our notation, the bags i for such a problem would be different starting locations in the problem search space³, and (the cumulants of) the $P_i(t)$ would be the runtime distributions of the algorithm when started at location i . $T(i, j)$ would most likely be a constant function describing the time cost of restarting the algorithm, To extend our theoretical results to such settings, it would hence be crucial to study *arbitrary* distributions $P_i(t)$ with constant $T(i, j) = T$, which would generalize Luby et al. [5] where $T(i, j) = 0$.

References

1. Stephens, D.W., Krebs, J.R.: Foraging Theory. Princeton University Press (1987)
2. Hofmeyr, S., Forrest, S.: Architecture for an artificial immune system. *Evolutionary Computation* **7**(1) (2000) 1289–1296
3. Hilker, M., Luther, K.: Artificial cell communication in distributed systems. In: AINA '08, Washington, DC, USA, IEEE Computer Society (2008) 1034–1041
4. Hoos, H., Stützle, T.: Stochastic Local Search: Foundations and Applications. Morgan Kaufmann (2005)
5. Luby, M., Sinclair, A., Zuckerman, D.: Optimal speedup of las vegas algorithms. *Inf. Process. Lett.* **47** (September 1993) 173–180
6. Janeway, C., Travers, P., Walport, M., Shlomchick, M.: Immunobiology. Garland Science (2005)
7. Blattman, J.N., Antia, R., Sourdiva, D.J., Wang, X., Kaech, S.M., Murali-Krishna, K., Altman, J.D., Ahmed, R.: Estimating the precursor frequency of naive antigen-specific CD8 T cells. *Journal of Experimental Medicine* **195**(5) (2002) 657–664
8. Westermann, J., Pabst, R.: Distribution of lymphocyte subsets and natural killer cells in the human body. *Clin. Investig* **70** (1992) 539–544
9. von Andrian, U.H.: Intravital microscopy of the peripheral lymph node microcirculation in mice. *Microcirculation* **3** (1996) 287–300
10. Wei, S.H., Parker, I., Miller, M.J., Cahalan, M.D.: A stochastic view of lymphocyte motility and trafficking within the lymph node. *Immunological Reviews* **195** (2003) 136–159
11. Glasser, M.L., Zucker, I.J.: Extended watson integrals for the cubic lattices. *PNAS* **74** (1977) 1800–1801
12. Weiss, G.H.: Asymptotic form for random walk survival probabilities on three-dimensional lattices with traps. *PNAS* **77**(8) (1980) 4391–4392
13. Soderberg, K.A., Payne, G.W., Sato, A., Medzhitov, R., Segal, S.S., Iwasaki, A.: Innate control of adaptive immunity via remodeling of lymph node feed arteriole. *PNAS* **102**(45) (2005) 16315–16320
14. Corless, R.M., Gonnet, G.H., Hare, D.E.G., Jeffrey, D.J., Knuth, D.E.: On the Lambert W function. *Advances in Computational Mathematics* **5** (1996) 329–359
15. Mempel, T.R., Henrickson, S.E., von Andrian, U.H.: T-cell priming by dendritic cells in lymph nodes occurs in three distinct phases. *Nature* **427** (2004) 154–159

³ Note that our definition does not require that the bags be disjoint regions in the problem space.