Coverage Maximization in Small World Network Under Bandwidth Constraint

Subrata Nandi
CSE, IIT Kharagpur, India
snandi@cse.iitkgp.ernet.in

Lutz Brusch
ZIH, TU Dresden, Germany
lutz@zhr.tu-dresden.de

Ajit Pal, Niloy Ganguly
CSE, IIT Kharagpur, India
{apal,niloy}@cse.iitkgp.ernet.in

ABSTRACT

This paper designs a coverage strategy which visits same number of distinct nodes much faster than the 1-random walk strategy by using same amount of bandwidth.

Categories and Subject Descriptors: C.2.0 [COMPUTER-COMMUNICATION NETWORKS]: General- *Data communications.*

General Terms: Design, Experimentation.

Keywords: Coverage strategy, random walk, proliferation, small world.

1. INTRODUCTION

This paper proposes a design strategy to maximize coverage[1] in large scale networks, under strict bandwidth[2] resource constraint. It is assumed that neither the nodes nor the spreading agents (query/information packets) maintain any additional information regarding node visits. Hence the naive coverage strategies are flooding and K-random walk.

Starting with multiple number of random walkers helps in reaching more number of nodes than a single random walker can reach during the same time. However, with multiple walkers redundant visits can occur both due to visit of the same walker to a particular node multiple number of times, as well as due to different walkers visiting the same node i.e. *mutual overlap*. With a single random walker only the first case can occur. Therefore, the probability of a redundant visit by a single walker in each step is minimum, which implies *final coverage* $\mathcal{FC}_{single} = \mathcal{FC}^{max}$. But the time $[B]$ required to obtain $\mathcal{FC}^{max}$ is too high to be acceptable for practical applications. On the other hand, although in general multiple walkers need less time latency to consume the provided bandwidth B, the final coverage $\mathcal{FC}_{multiple} < \mathcal{FC}^{max}$.

In this paper, we ask the very non-intuitive question that whether it is possible to introduce multiple walkers in the system in some ways so that both the objectives of obtaining maximum final coverage $\mathcal{FC}^{max}$ as well as consumption of bandwidth B at a much less time is achieved. To the best of our knowledge, the answer to this fundamental question has not yet been investigated; *even the question has not been posed.* In this paper taking cue from multiple random walker dynamics, we design a *proliferating random walker* based strategy, where walkers self-replicate avoiding mutual overlap, to achieve the maximum coverage $\mathcal{FC}^{max}$ in time $\mathcal{T} < B$. It may be that in many situations even $\mathcal{T} < B$ is not a practical latency limit. From our understanding of walker dynamics, we further propose strategy to maximize final coverage in such a situation. It is believed that many real world communication networks exhibit small world behavior. Hence in this paper we present the strategy on small world networks although it is applicable to other topologies.

[1]Coverage is the total number of distinct nodes visited.
[2]We assume throughout the paper that constraint in total bandwidth consumption is B and each packet forwarding consumes unit bandwidth.

SIGCOMM'08, August 17–22, 2008, Seattle, Washington, USA.
ACM 978-1-60558-175-0/08/08.

Table 1: Results on infinite Euclidean grid. Small world topology exhibits similar behavior.

Regimes :	I	II	III
C :	$d \times t^{d-1}$	$\frac{d}{2}[t \times \ln(K \times t^{1-\frac{d}{2}})]^{\frac{d}{2}-1}$	K
Time (t):	$---$	$-----------$	$--->$
Crossover time:	$O(\ln K)$ (I→II)		$O(K^{\frac{2}{d-2}})$ (II→III)

2. THE COVERAGE STRATEGY

Yuste[3] and Larralde[1] in their pioneering work modeled the random walker dynamics on an infinite d-dimensional Euclidean grid for K ($K \gg 1$) random walkers. *The first subsection designs the proliferation strategy for obtaining maximum final coverage* $\mathcal{FC}^{max}$ *in time* $\mathcal{T} < B$. *The second subsection designs the strategy while given time is* $T < \mathcal{T}$.

2.1 Bandwidth constraint

Random walker dynamics: The findings of [1] are summarized in Table 1, showing the presence of three distinct regimes. The *increase in coverage during one time step* is denoted by C. In regime I, all the walkers clutter together, hence mutual overlap probability is very high. As the walkers gradually move away from each other resulting in a decrease in the number of walkers at each node, the system enters regime II. However, still some amount of mutual overlap exists. In regime III, the walkers move sufficiently far apart from each other, such that mutual overlap is almost zero. As a result from time say, $t = t_K^{II-III}$ onwards i.e. *crossover time from regime II to III*, each walker behaves independently like a single random walker with non-overlapping exploration space.

From Table 1 we make some non intuitive **observations**: The coverage can be increased by using a larger value of K in regimes II and III, but the impact of which is significant only

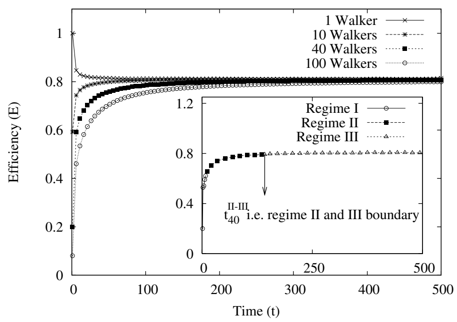

Figure 1: Plot of efficiency vs time varying walkers. The inset highlights the three regimes for 40 walkers

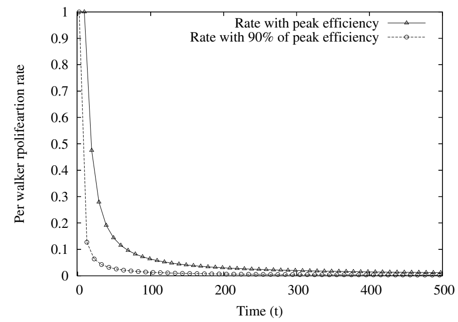

Figure 2: Plot of per walker proliferation rate vs time varying efficiency

in the regime III. The increase in K can however lengthen the duration of both regimes I and II. Let us denote E as the *efficiency during one time step*, which is defined as the ratio of the increase in coverage to the number of walkers used, during one time step i.e., $E = \frac{C}{K}$. The *peak efficiency* E^{max} is constant. It is achieved at $t = t_K^{II-III}$ i.e. the border of regimes II and III and is reached early if K is small. The crossover point between regimes II and III is an ideal point with maximum coverage and minimum overlap.

The above understanding yields the following **strategy**: *try to keep the system always in the crossover point of regimes II and III.* This can be achieved by starting with a small initial set of random walkers K_{init} and proliferating (multiplying) each walker at a suitable rate $\mathcal{P}_{rate}$ at each time step, so that the system always remains at the boundary.

Intuitive basis: The system with K walkers will crossover from regime II to III at time $t = t_K^{II-III}$. If no further walkers are introduced, the system will plunge deep into regime III, where the number of overlaps does not decrease further, but a lot of area remains unexplored. However, at this point, if a balanced number of walkers say, ΔK are introduced the system can just still stay in the border of the regimes. Let ΔK walkers be added at time t then the following condition must hold: $t + 1 = t_{K+\Delta K}^{II-III}$. Similarly if we want the system to be in the border of regimes II and III from the beginning, it should start with small number of walkers say, K_{init} where $t_{K_{init}}^{II-III} = 1$. We report the simulation results to verify the above mentioned dynamics in small world and then go on to estimate the *per walker proliferation rate* $\mathcal{P}_{rate}$.

Verification of the dynamics on small world: The small world network has been generated using Watts-Strogatz model[2] by adding 0.02% random links to a 4-dimensional torus Euclidean grid of size 10^8 nodes. Figure 1 shows the plot of efficiency E versus time t until time 500 units, for K =1, 10, 40 and 100 walkers. E is estimated as the mean of 10^5 simulation runs. From the plots we observe that E initially increases at a very fast rate (signifying regime I). Then gradually it slows down (signifying regime II) and finally settles to a constant value (signifying regime III), similar to the behavior shown in Table I. It is observed that the peak efficiency $E^{max} \approx 0.81$, is independent of K.

Estimating $\mathcal{P}_{rate}$ from simulation: $\mathcal{P}_{rate}$ ensures peak efficiency at each time step and thus keeping perennially the system in the border of regimes II and III. The proliferation rate can be estimated as $\mathcal{P}_{rate} = \frac{K_{t+1}|_{peak} - K_t|_{peak}}{K_t|_{peak}}$, where $K_t|_{peak}$ denotes the desired walker count required to ensure peak efficiency at time t. The $K_t|_{peak}$ is obtained from a synthesized function. At first for some chosen K's, $t = t^{II-III}|_K$ is obtained; then polyfit method is applied upon these $[K, t]$ pairs to derive the graph.

Results: The plot of $\mathcal{P}_{rate}$ versus time t is shown in Fig. 2, corresponding to the peak efficiency. The proliferation rate is a time varying function which decreases exponentially with time and asymptotically tends to zero. A little deeper probing into the dynamics reveals that at each time step the walkers can at best cover concentric hyper spheres with the increase in radius being almost constant over time. Hence the ratio $\mathcal{FC}_{t+1}/\mathcal{FC}_t$ decreases over time thus fueling the decrease in proliferation rate. For a given bandwidth B, $\mathcal{P}_{rate}$ produces $\mathcal{FC}^{max}$ in time $\mathcal{T}$. It is found that $\mathcal{T} \leq (0.1 \times B)$ for $B \geq 100$, implying 10 times speed-up. In the next section we deal with a more realistic situation when there is strict time constraint T such that $T < \mathcal{T}$.

2.2 Bandwidth and time constraint

This section designs coverage strategy which maximizes final coverage $\mathcal{FC}$ while taking time $T < \mathcal{T}$. The "excess bandwidth" which was originally spend between time T to $\mathcal{T}$ is now distributed in such a fashion so that a constant efficiency $E < E^{max}$ is ensured at each step.

The efficiency values E for various B and T are obtained through experimentations and numerical extrapolation. From the estimated efficiency E the per walker proliferation rate is obtained similar to section 2. The per walker proliferation rates corresponding to 90% of the peak efficiency is shown in Fig. 2. Note that the proliferation rates at both peak and 90% of the peak efficiency follow the same pattern.

Results: We compare the coverage obtained from the proposed strategy with equivalent strategies: constant α-proliferation and K-simple random walker, where α satisfies the condition $B = \Sigma(1 + \alpha)^T$ and $K = \frac{B}{T}$. It shows that for $B = 6000$ and $T = 69$, the proposed strategy covers almost 5% and 10% more nodes compared to the equivalent constant 0.097-proliferation and 87-walker strategy respectively. The performance improves more as T decreases.

3. CONCLUSION

The novelty of the work lies in the basic realization that one can achieve significant speed-up even without sacrificing efficiency and from there propose a proliferation strategy to realize such goal. Further experiments needed to be done to understand its behavior in various other topologies.

4. REFERENCES

[1] H. Larralde *et al*, "Territory covered by n diffusing particles," *Nature*, 355:423 – 426, Jan 2002.

[2] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks," *Nature*, 393:440 – 442, 1998.

[3] S. B. Yuste and L. Acedo. Number of distinct sites visited by n random walkers on an euclidean lattice," *Physical Review E*, 61(3):2340 – 2347, March 2000.