

# The price of evolution in incremental network design (the case of ring networks)\*

Saeideh Bakhshi, Constantine Dovrolis

College of Computing, Georgia Tech  
saeideh@gatech.edu, constantine@gatech.edu

**Abstract.** As it also happens in nature, technological networks typically evolve in an incremental manner, instead of being optimally designed. This evolutionary process is driven by changes in the underlying parameters and constraints (the “environment”) and it typically aims to minimize the modification cost after each change in the environment. In this paper, we first formulate the incremental network design approach and compare that with the more traditional optimized design approach in which the objective is to minimize the total network cost. We evaluate the cost overhead and evolvability of incremental design under two network expansion models (random and gradual), focusing on the simpler case of “ring” networks. We find that even though incremental design has some cost overhead, that overhead does not increase as the network grows. Also, it is less costly to evolve an existing network than to design it from scratch as long as the network expansion factor is less than a critical value.

## 1 Introduction

Complex technological systems, such as transportation and communication networks, manufacturing processes, microprocessors and computer operating systems, are rarely designed “from scratch.” Instead, they are often subject to an evolutionary process in which existing designs are incrementally modified every time a new phase or version of the system is needed. There are numerous examples and they span every engineering discipline. To mention one of them, consider a wide-area communication network that expands over time to reach new locations, increasing its capacity depending on the offered load, and occasionally providing new services.

It is remarkable that even though the optimized design of complex systems and networks has been studied in depth for several decades, the literature rarely considers that the design process is often incremental (or evolutionary). Instead, it is typically assumed that the system is designed *tabula rasa*, in a “clean-slate” manner. The corresponding design problems are typically formulated as

---

\* This research was supported by the National Science Foundation under Grant No. 0831848.

optimization problems with multiple constraints, none of which is imposed by an earlier design however. In the few previous studies that considered evolving networks (see Section 4), the focus was on the design algorithms and the corresponding optimization problems, rather than to examine the pros and cons of an incremental design relative the corresponding optimized design.

In this paper, we attempt a first investigation of the following fundamental question: *how does an incremental design compare to an optimized design, when both designs provide the same function?* Even though we believe that the previous question is relevant to the study of complex systems in general, we choose to focus on a rather narrow design problem and the engineering domain of our expertise, namely *the topological design of communication “ring” networks*. The ring topology is often used in practice mostly in metropolitan or regional networks [13]. In that context, the objective of the design process is to create a ring network that interconnects a given set of locations aiming to minimize a cost-related objective function. We further limit (admittedly in a very simplified manner) how the environment changes with time: the set of interconnected network locations expands by one or more nodes at each time step.

The previous ring topology design problem allows us to examine several interesting questions about incremental versus optimized designs in a precise and quantitative framework. How can we mathematically formulate the incremental design problem, and how is that formulation different from the more traditional optimized design problem? How does an incremental design compare, in terms of cost with the corresponding clean-slate design that provides the same function? How costly is it to modify an existing design, relative to the cost of re-designing the system from scratch? When is it better to abandon incremental changes on an existing system and start from scratch? What is the *“price of evolution”*, i.e., the cost overhead of an incremental design relative to the corresponding optimized design? What is the role of the pace at which the environment changes with time? Does the incremental design process perform better when the environment varies in certain non-random ways?

In Section 2, we present a formulation for the incremental design problem in which the objective is to *minimize the modification cost relative to the existing network*. We compare that formulation with an optimized design problem in which the objective is to minimize the total network cost. In Section 3, we derive expressions for the evolvability and cost overhead of the incremental design process. We also compare *random expansion* with *gradual expansion*. We review the related work in Section 4 and conclude in Section 5. In a follow-up paper we will expand this study in the case of general mesh networks.

## 2 Framework and Metrics

In this section, we present mathematical formulations for the clean-slate and incremental design problems. Even though these formulations are quite general, in the rest of the paper we apply them in the context of topology design for communication networks. As in any design problem, there is a desired function

(e.g., construct a communication network to connect a given set of locations), some design elements (e.g., routers, wide-area links), certain constraints (e.g., related to reliability or maximum propagation delay) as well as an objective (e.g., minimize the total cost of the required design elements). The design process aims to use the appropriate elements so that we achieve the desired function, while satisfying the constraints and meeting a given objective. It is often assumed that this process is conducted only once in an otherwise static environment. Here, we consider the case that the design takes place in a *dynamic environment*. In the context of communication networks, the network may gradually expand to new locations, the cost of design elements may fluctuate, or the constraints may become more stringent from time to time. We consider a discrete-time model and we refer to the  $k$ 'th time epoch as the  $k$ 'th *environment*. At a given environment  $k$ , all inputs of the design problem are known and constant.

How can we design a communication network in a dynamic environment? We identify two fundamentally different approaches. In the *clean-slate approach* we aim to *minimize in every environment the total cost of the network* subject to the given constraints. We refer to the resulting network in each environment as *optimized*. In the *incremental approach* we aim instead to *minimize the modification cost relative to the network of the previous environment*, again subject to the given constraints; we refer to the resulting network as *evolved*.

More rigorously, let  $\mathcal{N}(k)$  be the set of *acceptable networks* at environment  $k$ , i.e., networks that provide the desired function and meet the given constraints at environment  $k$ . The cost of a particular network  $N \in \mathcal{N}(k)$  is  $C(N)$ .  $C(N)$  is the sum of the costs of all design elements in  $N$ . We assume that there are no other costs associated with  $N$ ; for instance, there is no monetary cost to compute the design or to interconnect its elements.

In clean-slate design, the objective is to identify an acceptable network  $N_{opt}(k)$  from the set  $\mathcal{N}(k)$  that has the minimum cost  $C_{opt}(k)$  at environment  $k$ ,

$$C_{opt}(k) \equiv C(N_{opt}(k)), \quad N_{opt}(k) \equiv \arg \min_{N \in \mathcal{N}(k)} C(N). \quad (1)$$

We refer to  $N_{opt}(k)$  as the *optimized* network at environment  $k$ . If the optimized network is not unique, we break ties with secondary objectives (for instance, minimize the total number of links). Most network design problems are computationally intractable (NP-hard), and so they are often solved heuristically, approximating the previous optimization objective. This is what we also do in Section 3. For this reason, we do not refer to  $N_{opt}(k)$  as *optimal* but as *optimized*. The former would be the actual solution to the previous problem if we could compute it; the latter is the best solution we can compute given a certain design heuristic.

In the incremental design approach, on the other hand, we design the new network  $N_{evo}(k)$  based on the network  $N_{evo}(k-1)$  from the previous environment  $k-1$ . We refer to the former as the *evolved* network at environment  $k$ . The objective of the incremental design process is to identify an acceptable network  $N(k) \in \mathcal{N}(k)$  that minimizes the *modification cost*  $C_{mod}(N_{evo}(k-1); N(k))$

between networks  $N_{evo}(k-1)$  and  $N(k)$ . For simplicity, we denote the previous modification cost as  $C_{mod}(k)$ .

The modification cost  $C_{mod}(k)$  is defined as the cost of new design elements that are needed in  $N(k)$  but are not present in  $N_{evo}(k-1)$ . Formally,  $C_{mod}(k)$  is the cost of the design elements in the set  $S_{mod}(k)$ , where

$$S_{mod}(k) = N(k) \setminus [N_{evo}(k-1) \cup I(k-1)] \quad (2)$$

abusing the notation  $N(k)$  to also refer to the set of design elements in the network  $N(k)$ .

With the previous definitions, we can now formulate the incremental design process as:

$$C_{evo}(k) \equiv C(N_{evo}(k)), \quad N_{evo}(k) \equiv \arg \min_{N \in \mathcal{N}(k)} C_{mod}(k) \quad (3)$$

The evolved network  $N_{evo}(k)$  may not be unique in general. Ties are broken by considering a secondary objective: if two networks minimize the modification cost, select the network with the minimum total cost.<sup>1</sup> As in the case of optimized design, we compute the solution of the incremental design problem with a heuristic described in Section 3.

The cost of the evolved network can be expressed recursively as ( $k \geq 1$ ):

$$C_{evo}(k) = C_{evo}(k-1) + C_{mod}(k) \quad (4)$$

We assume that the initial evolved network (and its cost  $C_{evo}(0)$ ) is known. So, the cost of the evolved network at environment  $k$  has increased by the modification cost (new elements that are purchased at  $k$ ).

Expanding (4), we can write the cost of the evolved network as:

$$C_{evo}(k) = C_{evo}(0) + \sum_{i=1}^k C_{mod}(i) \quad (5)$$

Thus, the cost of the evolved network at environment  $k$  is the cost of the initial network plus the cost of all design elements that were purchased in the last  $k$  environments.

## 2.1 Metrics

We now introduce four metrics to compare a sequence of optimized and evolved networks. We also introduce two specific models of dynamic environment we consider in this paper, and quantify the rate at which the environment changes with time.

---

<sup>1</sup> In our computational experiments, two modification costs are rarely equal because link costs are based on distance and they are real numbers.

First, the *cost overhead*  $v(k)$  of the evolved design  $N_{evo}(k)$  relative to the corresponding optimized design  $N_{opt}(k)$  at environment  $k$  is:

$$v(k) = \frac{C_{evo}(k)}{C_{opt}(k)} - 1 \geq 0 \quad (6)$$

where the inequality is expected from the definition of  $C_{opt}(k)$ .<sup>2</sup> What is more important however is whether the cost overhead of the incremental design process gradually increases, i.e., whether the evolved networks become increasingly more expensive compared to the corresponding optimized networks. If that is the case, the incremental design process would diverge over the long-term towards extremely inefficient designs.

Second, the *evolvability*  $e(k)$  is defined as:

$$e(k) = 1 - \frac{C_{mod}(k)}{C_{opt}(k)} \leq 1. \quad (7)$$

The evolvability represents the cost of modifying the evolved network from environment  $k-1$  to  $k$ , relative to the cost of redesigning the network “from scratch” at time  $k$ . High evolvability, close to 1, means that it is much less expensive to modify the existing network than to re-design a new network. On the other hand, when the evolvability becomes zero or negative, it is beneficial to stop the incremental design process and design a new optimized network, i.e., clean-slate design “beats” evolution in that case.

The *topological similarity*  $t(k)$  between the optimized  $N_{opt}(k)$  and evolved  $N_{evo}(k)$  networks is defined as the *Jaccard similarity coefficient* of the two corresponding adjacency matrices. In other words,  $t(k)$  is the fraction of distinct links in either network that are present in both networks. Even though the two network design approaches we consider are different in terms of objective and design method, we are interested to know how different the resulting networks are, structure-wise, over time.

## 2.2 Expansion models

We consider a specific way in which the environment changes with time: *expansion*. Specifically, the set of locations that the network has to interconnect at any time  $k$  is increasing with  $k$ . This is probably the most natural way the environment can change with time in the context of communication networks.

In the simplest form of expansion, the network size increases by only one node at each environment; we refer to this as *single-node expansion*.

We also consider a *multi-node expansion* scenario in which the network size increases once by a multiplicative factor  $\rho$ , which we refer to as *expansion factor*.

---

<sup>2</sup> The reader should note that the cost overhead metric is different than the well-known *approximation ratio*. The latter examines the worse-case solution produced by an algorithm relative to the optimal solution of a given problem. The cost overhead compares the solutions (costs) of two algorithms that aim to minimize two *different* optimization objectives.

Specifically, if the network size increases from  $n$  nodes to  $n + m$  nodes, the expansion factor  $\rho$  is

$$\rho = \frac{n + m}{n} \geq 1. \quad (8)$$

We also compare two expansion models: *random* and *gradual*. Suppose that the set of all possible locations is  $\mathcal{L}$  and that the network expands at time  $k$  to  $X$  new locations. In random expansion, the  $X$  new locations are selected randomly from  $\mathcal{L}$ . In gradual expansion, we select iteratively each of the  $X$  new locations from  $\mathcal{L}$  so that it is the closest location to either any of the existing nodes in network  $N(k - 1)$  or to any of the new locations we have just added in the network.

The random and gradual expansion models represent two significantly different models in which the environment changes with time. In random expansion, the new locations can be anywhere and so it may be costly for the incremental design to adjust the previous network with only minor modifications. In gradual expansion, the environment changes in an “evolution-friendly” manner because the new locations are as close as possible to the existing network.

### 3 Ring Networks

In this section, we compare the optimized and incremental design approaches in the context of a ring topology. Ring networks are widely used mostly in metropolitan-area networks, as they are robust to single-node failures (two node-disjoint paths exist between any pair of nodes) and they are typically less costly than mesh networks [9]. For our purposes, the ring topology has two additional features. First, we can use an existing software package (Concorde) [8] that computes excellent approximations to the optimal ring design problem. Second, we can leverage existing analytical results to derive asymptotic expressions for the evolvability and cost overhead of incremental ring design under random and gradual expansion.

#### 3.1 Optimized versus incremental ring design

We assume that all potential nodes of the expanding ring network are located in a bounded plane region. Further, we assume that the cost of a network is equal to the sum of its link costs, and the cost of each link is proportional to its length. So, the minimum-cost ring design problem is equivalent to the NP-Hard Traveling Salesman Problem (TSP) [1]. We rely on a TSP-solving software package called Concorde [8], which is based on a branch-and-cut algorithm. Concorde’s TSP solver has computed the *optimal* solutions to 106 of the 110 *TSPLIB* problem instances; the largest of them has 15,112 nodes.

Further, in the case of ring design we can use a simple asymptotic expression for the length of the optimal TSP tour. Specifically, Beardwood et al. proved that “*the length of the shortest closed path through  $n$  points in a bounded plane region of area  $A$  is almost always asymptotically proportional to  $\sqrt{A \times n}$  for*

large  $n$ " [3]. In our context, the length of the TSP tour is equal to the cost of the optimized ring.

In the case of random expansion, the  $n$  nodes of the ring can be anywhere in the given region and so the area  $A$  does not depend on  $n$ . Thus, the cost of the optimized ring increases as


$$C_{opt}^{rnd}(n) \sim \sqrt{A \times n} \sim \sqrt{n} \quad (9)$$

where the notation  $x \sim f(n)$  means that  $x$  tends to become proportional to  $f(n)$  as  $n$  increases.

In the case of gradual expansion, the area in which the  $n$  ring nodes are located increases with  $n$ . If we assume that all possible nodes are uniformly distributed with point density  $\sigma_{grd}$  in a bounded plane region, then the area  $A$  in which  $n$  ring nodes are located increases proportionally with  $n$  ( $n = A \sigma_{grd}$ ). Thus,

$$C_{opt}^{grd}(n) \sim \sqrt{\sigma_{grd} \times n^2} \sim n. \quad (10)$$

In the incremental ring design process, we can compute the minimum modifica-



**Fig. 1.** Connecting new nodes to an existing ring: (a) single-node expansion, (b) multi-node expansion (we show the order in which nodes are connected).

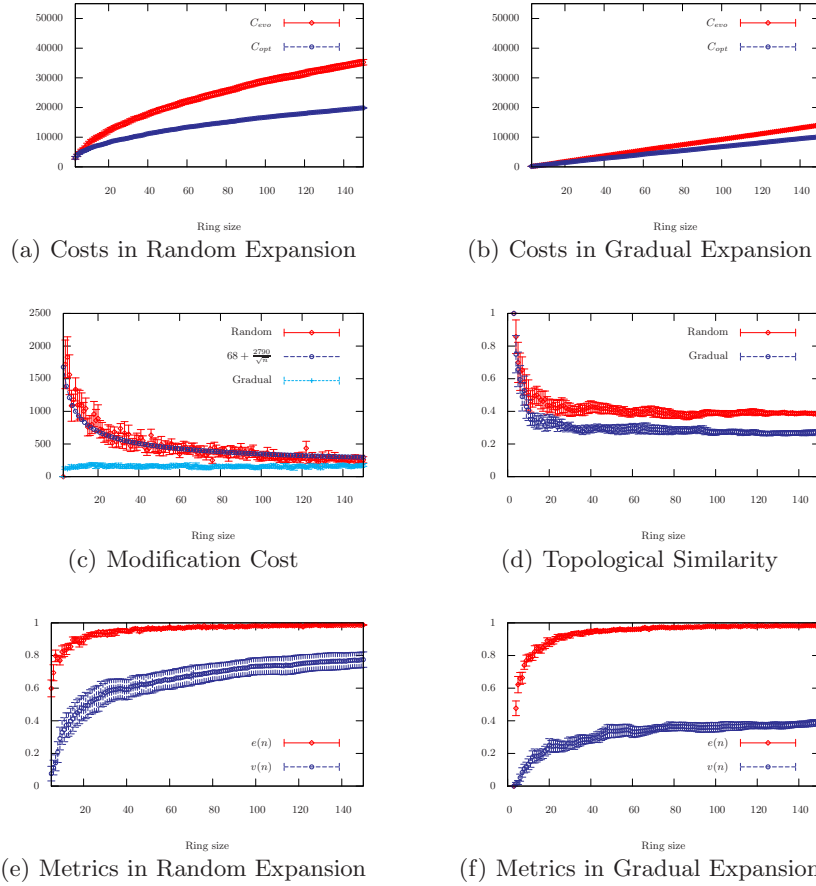
tion cost under single-node expansion. Suppose that the existing ring  $N_{evo}(k-1)$  has size  $n$  and we add a single extra node  $z$  at time  $k$ . The minimum modification cost will result if we connect  $z$  to two adjacent nodes  $x$  and  $y$  of  $N_{evo}(k-1)$ , such that

$$C_{mod}(k) = \min_{(x,y) \in N_{evo}(k-1)} (||z - x|| + ||z - y||) \quad (11)$$

and then removing the edge  $(x, y)$ . This process is illustrated in Figure 1(a). Note that there is no other way to connect  $z$  to  $N_{evo}(k-1)$  so that the resulting network is still a ring but with lower modification cost.

In the case of multi-node expansion, we use an iterative heuristic that aims to minimize the modification cost. Suppose that the existing ring  $N_{evo}(k-1)$  has size  $n$  and we add a set  $Z$  of more than one new nodes at time  $k$ . In each iteration, we select the node  $z$  from  $Z$  that minimizes the expression (11), connect  $z$  to the existing ring as we do in the case of single-node expansion, and then

move  $z$  from  $Z$  to the set of nodes in  $N_{evo}(k-1)$ . This process is illustrated in Figure 1(b). Note that this greedy heuristic may be sub-optimal in minimizing the modification cost  $C_{mod}(k)$  - an exhaustive search for the optimal solution however can be prohibitively slow.



**Fig. 2.** Results for ring network under random and gradual single-node expansion.

### 3.2 Random single-node expansion

To simplify the notation, instead of referring to an environment  $k$ , we simply refer to the ring size  $n$ . So, instead of writing  $C_{mod}(k)$  we write  $C_{mod}(n)$ , which refers to the modification cost when the ring expands to  $n$  nodes.

To compute the modification cost under random expansion, we rely on a result for the nearest neighbor problem: given a set  $S$  of  $n$  points and a new point  $z$  find the closest neighbor of  $z$  in  $S$ . When  $n$  points reside in a two-dimensional region of size  $A$  with point density  $\sigma_{rnd}$ , the expected value of the nearest neighbor distance is  $1/\sqrt{\sigma_{rnd}} = 1/\sqrt{n/A}$  [2]. The modification cost in Equation (11) can be approximated as twice the distance between the new node and the nearest node in the ring. The expected value of the latter is  $1/\sqrt{n/A}$  because the  $n$  ring nodes are randomly placed in the region  $A$ . So, based on the previous approximation, the modification cost  $C_{mod}(n)$  also scales as the mean nearest neighbor distance,

$$C_{mod}^{rnd}(n) \sim \frac{1}{\sqrt{n}}. \quad (12)$$

We have confirmed the validity of the previous approximation with computational results (see Figure 2(c)).

Using Equation (5), we see that the cost of the evolved network has the same scaling behavior as the cost of the optimized network,

$$C_{evo}^{rnd}(n) = \sum_{i=2}^n C_{mod}^{rnd}(i) \sim \sqrt{n}. \quad (13)$$

We can now derive asymptotic expressions for the evolvability and cost overhead under random expansion. From (9) and (12), it follows that

$$1 - e^{rnd}(n) \sim \frac{1}{n}. \quad (14)$$

Thus, the evolvability under single-node expansion converges to one, i.e., for large rings, the modification cost is practically zero compared to the cost of designing a new optimized ring. Also, we expect that the cost overhead will be practically constant for large values of  $n$  because both the evolved and optimized network costs scale as  $\sqrt{n}$ ,

$$v^{rnd}(n) \sim \text{constant}. \quad (15)$$

Thus, the evolved ring does *not* become increasingly more expensive relative to the optimized ring under random single-node expansion. The exact value of the cost overhead depends on the placement of the nodes, the order in which they are added to the network, and the initial ring we start from.

We have confirmed these asymptotic expressions with computational experiments in which the optimized ring is designed using Concorde and the evolved ring is designed based on Equation 12. Figure 2(a) shows the optimized and evolved network costs (with 90% confidence intervals for the empirical results after 20 runs), Figure 2(c) shows the modification cost, while Figure 2(e) shows the evolvability and cost overhead. In the last graph, note that the cost overhead appears to increase with the ring size - we have confirmed that for larger ring sizes (larger than approximately 200 nodes) the cost overhead converges to about 0.8.

The topological similarity between the corresponding optimized and evolved rings is shown in Figure 2(d). It is interesting that the two networks share only about 30-40% of their links, when  $n > 20$ . Thus, the two network topologies become (and stay) significantly different during the expansion process (even though they started from the same three-node ring). Also, the topological similarity is slightly higher in random expansion compared to gradual expansion.

### 3.3 Gradual single-node expansion

In this case, a new node is selected among all potential locations as the closest location to any existing ring node. We can rely again on the nearest neighbor problem to estimate the modification cost. Suppose that all potential new locations are uniformly distributed in the given area with density  $\sigma_{grad}$ . The expected value of the distance between a node in the current ring (of size  $n$ ) and the nearest potential location is  $1/\sqrt{\sigma_{grad}}$ , which does not depend on  $n$ . The modification cost can be again approximated by twice the previous distance, and so it should not increase with  $n$ , at least for large rings,

$$C_{mod}^{grad}(n) \sim \text{constant}. \quad (16)$$

Using Equation (5), we see that the cost of the evolved network under gradual expansion scales linearly with  $n$ ,

$$C_{evo}^{grad}(n) = \sum_{i=2}^n C_{mod}^{grad}(i) \sim n. \quad (17)$$

So, the evolvability under gradual expansion scales as in the case of random expansion

$$1 - e^{grad}(n) \sim \frac{1}{n}. \quad (18)$$

The cost overhead does not increase with  $n$  for large values of  $n$ , as in the case of random expansion,

$$v^{grad}(n) \sim \text{constant} \quad (19)$$

based on Equations (17) and (10).

The difference between random and gradual single-node expansion becomes evident in the computational results (see Figure 2). For the same ring size, the evolved and optimized costs are lower under gradual expansion, the modification cost is practically constant (and lower than under random expansion), the evolvability is not significantly different between the two expansion models, while the cost overhead is significantly lower under gradual compared to random expansion. In other words, when the ring expands in a gradual manner, we expect that the evolved network will be closer, in terms of cost, to the optimized network compared to random expansion.

### 3.4 Effect of expansion factor $\rho$

Suppose that we add  $m$  new nodes to a ring of size  $n$  so that the resulting network is a ring with  $n + m$  nodes. The expansion factor is  $\rho = \frac{n+m}{n} > 1$ . What can we expect about the evolvability and cost overhead of the incremental design process as functions of  $\rho$ ?

If  $m \ll n$  (i.e.,  $\rho$  is close to one), we can rely on the following simple approximation. Under random expansion, the modification cost  $C_{mod}^{rnd}(n; m)$  when we add  $m$  new nodes will be approximately  $m$  times larger than the modification cost  $C_{mod}^{rnd}(n)$  when we add a single new node at a ring of size  $n$ . So,

$$C_{mod}^{rnd}(n; m) \approx m C_{mod}^{rnd}(n) \sim m/\sqrt{n}. \quad (20)$$

On the other hand, the cost of an optimal ring with  $n + m$  nodes will be

$$C_{opt}^{rnd}(n; m) \sim \sqrt{n + m}. \quad (21)$$

Thus, the evolvability under random expansion scales as

$$1 - e^{rnd}(n; m) = \frac{C_{mod}^{rnd}(n; m)}{C_{opt}^{rnd}(n; m)} \sim \frac{m}{\sqrt{n(n + m)}} \quad (22)$$

that can be written as

$$1 - e^{rnd}(n; m) \sim \frac{\rho - 1}{\sqrt{\rho}}. \quad (23)$$

As  $\rho$  increases the evolvability decreases and there is a critical expansion factor value  $\hat{\rho}$  at which the evolvability becomes zero. For larger expansions than  $\hat{\rho}$ , it is better to abandon the existing network and redesign the ring in a clean-slate manner.

In the case of gradual expansion, the modification cost  $C_{mod}^{grd}(n; m)$  will again be roughly  $m$  times larger than the modification cost  $C_{mod}^{grd}(n)$  when we add only one node,

$$C_{mod}^{grd}(n; m) \approx m C_{mod}^{grd}(n) \sim m \times \text{constant} \quad (24)$$

while the cost of an optimal ring with  $n + m$  nodes is

$$C_{opt}^{grd}(n; m) \sim n + m. \quad (25)$$

Thus, the evolvability under gradual expansion scales as

$$1 - e^{grd}(n; m) = \frac{C_{mod}^{grd}(n; m)}{C_{opt}^{grd}(n; m)} \sim \frac{m}{n + m} \quad (26)$$

that can be written as

$$e^{grd}(n; m) \sim \frac{1}{\rho}. \quad (27)$$

Again, as  $\rho$  increases the evolvability decreases, but more slowly than under random expansion. It is also important that a critical expansion factor at which

the evolvability becomes zero may *not* exist under gradual expansion. Equation (27) was derived assuming that  $\rho$  is close to one, and so we cannot rely on that expression to prove that the evolvability is always positive. Computational results, however, indicate that this may be the case under gradual expansion (see Figure 3(a)).

To derive the cost overhead under multi-node expansion, we have to make two additional assumptions. First, the evolved network at size  $n$  is the same with the optimized network at size  $n$  (i.e.,  $C_{evo}(n) = C_{opt}(n)$ ). Second, because  $m \ll n$ , the cost of the optimized network at size  $n$  is approximately equal to the cost of the optimized network at size  $n + m$  (i.e.,  $C_{opt}(n) \approx C_{opt}(n + m)$ ). Then, under both random and gradual expansion, we can write that

$$C_{evo}(n + m) = C_{evo}(n) + C_{mod}(n; m) \approx C_{opt}(n + m) + m C_{mod}(n) \quad (28)$$

and so the cost overhead is

$$v(\rho) = \frac{C_{evo}(n + m)}{C_{opt}(n + m)} - 1 \approx m \frac{C_{mod}(n)}{C_{opt}(n)}. \quad (29)$$

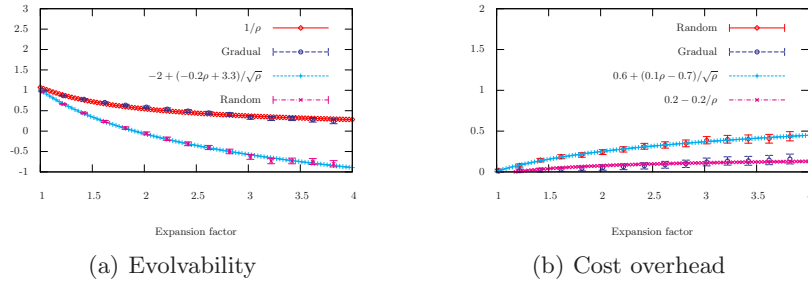
Thus, under random expansion, the cost overhead scales as

$$v^{rnd}(\rho) \sim \frac{\rho - 1}{\sqrt{\rho}} \quad (30)$$

while under gradual expansion it scales as

$$v^{grad}(\rho) \sim 1 - \frac{1}{\rho}. \quad (31)$$

The previous scaling expressions are derived assuming that  $\rho$  is close to one, but computational results confirm that they are quite accurate when  $\rho$  is as high as four (see Figure 3(b)).



**Fig. 3.** Effect of expansion factor under random and gradual expansion.

We conducted computational experiments in which an initial ring (of size 50) is increased, under random or gradual multi-node expansion, by different factors

$\rho$ . Figure 3.4 shows the evolvability and cost overhead as functions of  $\rho$ . The previous scaling expressions, which are derived assuming that  $n$  is large and  $\rho$  is close to one, are actually in close agreement with the numerical results when  $\rho$  is as high as four. Under random expansion, the critical expansion factor is  $\hat{\rho} \approx 2$ , while the corresponding cost overhead is approximately 25%, meaning that even though it is then beneficial to re-design the network from scratch, the cost overhead of the evolved network is still quite low. Under gradual expansion, the evolvability does not become negative, at least in the given range of  $\rho$ , and the cost overhead is significantly lower than under random expansion.

## 4 Related work

The topology design literature is extensive both in the domain of computer networks and in theoretical computer science, and it is well covered in a recent book by Pioro and Medhi [13]. The vast majority of that literature, however, focuses on optimized network design. Only few studies have focused on incremental network design (also referred to as “multi-period design”), and none of them, to the extent of our knowledge, have focused on a comparison between incremental and optimized design.

Specifically, there are some studies that focus on incremental network design [11, 15, 10, 7, 4, 16, 6, 12, 17, 18] (see also section 11.2 of [13]). Those works mostly propose algorithms and optimization frameworks for incremental network design under a wide range of different constraints and objectives. Some of them consider topology design while others consider capacity expansion coupled with routing changes, and some of them consider reliability constraints while others consider limited budget constraints. None of them, however, is significantly relevant to our study because they do not compare incremental designs with the corresponding optimized designs, and they do not consider different expansion models (e.g., random versus gradual or single-node versus multi-node).

Chiang and Yang have focused on various “X-ities”, such as evolvability, scalability, reliability, or adaptability in the context of computer networks [5]. They present an analytic framework to capture the notions of evolvability and scalability. They also present an Evolvable Network Design (END) Tool using dynamic programming methods to design the multi-phase deployment of a network so that early phase designs are more evolvable in later stages.

A quite different, but still relevant, study by Tero et al. [14] compared the Tokyo rail system (as an example of an optimally designed transportation network) with a natural network formed by the slime mold *Physarum polycephalum*. The slime mold was allowed to grow on a rectangular map of the city of Tokyo; the map contained food on the locations at which the Tokyo rail system has stations. The slime mold network grew in an incremental manner, without any centralized control or “intelligence”, and it gradually covered and interconnected all food locations while refining its connectivity over time. The authors compared the two networks in terms of efficiency, fault tolerance, and cost and found that they are actually quite similar! This novel experiment implies that even a sim-

ple and incremental design process may be able to produce a network that has the cost efficiency and reliability properties that we usually only expect from optimized networks.

## 5 Conclusion

We can now return to the questions that were asked in the introduction and summarize our main findings. The following conclusions are supported by asymptotic scaling expressions as well as computational results for rings.

1. We formulated the incremental ring design process as an optimization problem that aims to minimize the modification cost relative to the previous network. We also identified and compared certain expansion models (random versus gradual, and single-node versus multi-node).
2. Under single-node expansion (random and gradual), even though an evolved network has higher cost than the corresponding optimized network, both costs scale similarly ( $\sqrt{N}$ ). As a result the the cost overhead of evolved network does not increase as the network grows.
3. Under single-node (random or gradual) expansion, it is less costly to follow the incremental design approach than to re-design the network from scratch. The evolvability under basic expansion converges to one as the network grows.
4. Under multi-node and random expansion, there is a critical value  $\hat{\rho}$  of the expansion factor beyond which it is less costly to abandon the existing network and re-design the network from scratch. In the gradual expansion, the evolvability scales as  $1/\rho$ .
5. Under gradual expansion, the evolvability is higher and the cost overhead is lower than under random expansion. The model of gradual expansion represents a more “evolution-friendly” dynamic environment than random expansion.

In a follow-up paper we will extend this study to capture general, mesh networks. In that case, networks are designed with additional reliability and performance constraints (e.g., a minimum number of node-disjoint paths between any pair of nodes, or a maximum propagation delay between any two nodes). We will also examine other dynamic environment models, such as iterated multi-node expansion, models in which the traffic loads and link capacities change with time as well, as well as more elaborate economic models that involve discounting or dynamic costs. It would also be interesting to see this quantitative framework and comparisons between evolved and optimized designs applied in other problems and technological domains.

## References

1. D. Applegate, R. Bixby, V. Chvatal, and W. Cook. *The Traveling Salesman Problem: (A Computational Study)*. Princeton University Press, 2006.
2. S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *J. ACM*, 45:891–923, 1998.
3. J. Beardwood, J. Halton, and J. Hammersley. The shortest path through many points. *Mathematical Proceedings of the Cambridge Philosophical Society*, (55):299–327, 1959.
4. P. Carvalho, L. Ferreira, F. Lobo, and L. Barruncho. Optimal distribution network expansion planning under uncertainty by evolutionary decision convergence. *International Journal of Electrical Power & Energy Systems*, 20(2):125–129, 1998.
5. M. Chiang and M. Yang. Towards Network X-ities From a Topological Point of View: Evolvability and Scalability. *Proc., Allerton Conf. on Comm., Control, and Computing*, 2004.
6. N. Geary, A. Antonopoulos, E. Drakopoulos, and J. O’Reilly. Analysis Of Optimisation Issues In Multi-Period DWDM Network Planning. In *IEEE INFOCOM*, 2001.
7. S. Gopal and Krishna Jain. On Network Augmentation. *IEEE Transactions on Reliability*, 35(5):541–543, 1986.
8. M. Hahsler and K. Hornik. TSP Infrastructure for the Traveling Salesperson Problem. *IEEE/ACM Transactions on Networking*, 23(2), December 2007.
9. M. Herzog, M. Maier, and M. Reisslein. Metropolitan area packet-switched wdm networks: A survey on ring systems. *Communications Surveys & Tutorials, IEEE*, 6(2):2–20, 2004.
10. H. Lee and D. Dooly. Heuristic algorithms for the fiber optic network expansion problem. *Telecommunication Systems*, 7:355–378, 1997.
11. A. Meyerson, K. Munagala, and S. Plotkin. Designing Networks Incrementally. *IEEE FOCS*, 2001.
12. M. Pickavet and P. Demeester. Long-term planning of WDM networks: A comparison between single-period and multi-period techniques. *Photonic Network Communications*, 1(4):331–346, 1999.
13. M. Pioro and D. Medhi. *Routing, Flow and Capacity Design in Communication and Computer Networks*. The Morgan Kaufmann Series in Networking, 2004.
14. A. Tero, S. Takagi, T. Saigusa, K. Ito, D. Bebbber, M. Fricker, K. Yumiki, R. Kobayashi, and T. Nakagaki. Rules for Biologically Inspired Adaptive Network Design. *Science*, 327(5964):439–442, 2010.
15. K. Vajanapoom and D. Tipper. Risk based incremental survivable network design. 2007.
16. T. Wu, R. Cardwell, and M. Broyden. A multi-period design model for survivable network architecture selection for SONET interoffice networks. *IEEE Trans. on Reliability*, 40:417–427, October 1991.
17. B. Yaged. Minimum Cost Routing for Dynamic Network Models. *Networks*, 3:193–224, 1973.
18. N. Zadeh. On Building Minimum Cost Communication Networks over time. *Networks*, 4:19–34, 1974.