

ClouDiA: a deployment advisor for public clouds

Tao Zou · Ronan Le Bras · Marcos Vaz Salles ·
Alan Demers · Johannes Gehrke

Received: 1 February 2014 / Revised: 11 October 2014 / Accepted: 18 December 2014 / Published online: 24 January 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract An increasing number of distributed data-driven applications are moving into shared public clouds. By sharing resources and operating at scale, public clouds promise higher utilization and lower costs than private clusters. To achieve high utilization, however, cloud providers inevitably allocate virtual machine instances non-contiguously; i.e., instances of a given application may end-up in physically distant machines in the cloud. This allocation strategy can lead to large differences in average latency between instances. For a large class of applications, this difference can result in significant performance degradation, unless care is taken in how application components are mapped to instances. In this paper, we propose ClouDiA, a general deployment advisor that selects application node deployments minimizing either (i) the largest latency between application nodes, or (ii) the longest critical path among all application nodes. ClouDiA employs a number of algorithmic techniques, including mixed-integer programming and constraint programming techniques, to efficiently search the space of possible mappings of application nodes to instances. Through experiments with synthetic and real applications in Amazon EC2, we show that mean latency is a robust metric to model

communication cost in these applications and that our search techniques yield a 15–55 % reduction in time-to-solution or service response time, without any need for modifying application code.

Keywords Embedding · Cloud computing · Parallel frameworks · Database optimizations

1 Introduction

With advances in data center and virtualization technology, more and more distributed data-driven applications, such as high-performance computing (HPC) applications [24, 34, 51, 72], web services and portals [28, 44, 52, 60], and even search engines [6, 8], are moving into public clouds [4]. Public clouds represent a valuable platform for tenants due to their incremental scalability, agility, and reliability. Nevertheless, the most fundamental advantage of using public clouds is cost-effectiveness. Cloud providers manage their infrastructure at scale and obtain higher utilization by combining resource usages from multiple tenants over time, leading to cost reductions unattainable by dedicated private clusters.

Typical cloud providers adopt a pay-as-you-go pricing model in which tenants can allocate and terminate virtual machine instances at any time and pay only for the machine hours they use [2, 49, 66]. Giving tenants such freedom in allocating and terminating instances, public clouds face new challenges in choosing the placement of instances on physical machines. First, they must solve this problem at scale, and at the same time take into account different tenant needs regarding latency, bandwidth, or reliability. Second, even if a fixed goal such as minimizing latency is given, the placement strategy still needs to take into consideration the possibility

T. Zou (✉) · R. Le Bras · A. Demers · J. Gehrke
Cornell University, Ithaca, NY, USA
e-mail: taozou@cs.cornell.edu

R. Le Bras
e-mail: lebras@cs.cornell.edu

A. Demers
e-mail: ademers@cs.cornell.edu

J. Gehrke
e-mail: johannes@cs.cornell.edu

M. V. Salles
University of Copenhagen, Copenhagen, Denmark
e-mail: vmarcos@diku.dk

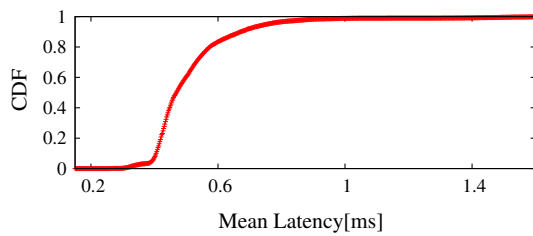


Fig. 1 Latency heterogeneity in EC2

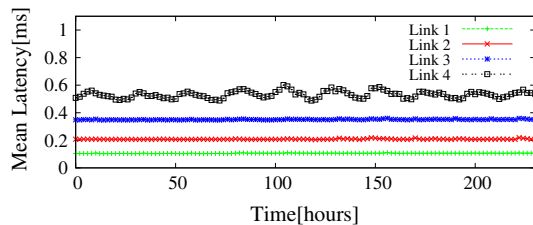


Fig. 2 Mean latency stability in EC2

of future instance allocations and terminations, not only for the current tenant, but also for other tenants who are sharing the resources.

Given these difficulties, public cloud service providers do not currently expose instance placement or network topology information to cloud tenants.¹ While these API restrictions ease application deployment, they may cost significantly in performance, especially for latency-sensitive applications. Unlike bandwidth, which can be quantified in the SLA as a single number, network latency depends on message sizes and the communication pattern, both of which vary from one application to another.

In the absence of placement constraints by cloud tenants, cloud providers are free to assign instances to physical resources non-contiguously; i.e., instances allocated to a given application may end-up in physically distant machines. This leads to heterogeneous network connectivity between instances: Some pairs of instances are better connected than other pairs in terms of latency, loss rate, or bandwidth. Figure 1 illustrates this effect. We present the CDF of the mean pairwise end-to-end latencies among 100 Amazon EC2 large instances (m1.large) in the US East region, obtained by TCP round-trip times of 1 KB messages. Around 10 % of the instance pairs exhibit latency above 0.7 ms, while the bottom 10 % are below 0.4 ms. This heterogeneity in network latencies can greatly increase the response time of distributed, latency-sensitive applications. Figure 2 plots the mean latencies of four representative links over a 10-day experiment,

¹ The only exception we know of is the notion of *cluster placement groups* in Amazon EC2 cluster instances. However, these cluster instances are much more costly than other types of instances, and only a limited number of instances can be allocated within a cluster placement group.

with latency measurements averaged every 2 h. The observed stability of mean latencies suggests that applications may obtain better performance by selecting “good” links for communication. In “Appendix 3,” we show that this intuition is not restricted to Amazon web services: Similar observations of latency heterogeneity and mean latency stability can be made in other main public cloud service providers, namely Google Compute Engine and Rackspace Cloud Server.

This paper examines how developers can carefully tune the deployment of their distributed applications in public clouds. At a high level, we make two important observations: (i) If we carefully choose the mapping from nodes (components) of distributed applications to instances, we can potentially prevent badly interconnected pairs of instances from communicating with each other; (ii) if we over-allocate instances and terminate instances with bad connectivity, we can potentially improve application response times. These two observations motivate our general approach: A cloud tenant has a target number of components to deploy onto x virtual machines in the cloud. In our approach, she allocates x instances plus a small number of additional instances (say $x/10$). She then carefully selects which of these $1.1 \cdot x$ instances to use and how to map her x application components to these selected virtual machines. She then terminates the $x/10$ over-allocated instances.

Our general approach could also be directly adopted by a cloud provider—potentially at a price differential—but the provider would need to widen its API to include latency-sensitivity information. Since no cloud provider currently allows this, we take the point of view of the cloud tenant, by whom our techniques are immediately deployable.

1.1 Contributions of this paper

In this paper, we introduce the problem of deployment advice in the cloud and instantiate a concrete deployment advisor called ClouDiA (*Cloud Deployment Advisor*).

1. ClouDiA works for two large classes of data-driven applications. The first class, which contains many HPC applications, is sensitive to the worst-link latency, as this latency can significantly affect total time-to-solution in a variety of scientific applications [1, 12, 22, 35]. The second class, represented by search engines as well as web services and portals, is sensitive to the longest path between application nodes, as this cost models the network links with the highest potential impact on application response time. ClouDiA takes as input an application communication graph and an optimization objective, automatically allocates instances, measures latencies, and finally outputs an optimized node deployment plan (Sect. 2). To the best of our knowledge, our methodology is the first to address deployment tuning for latency-sensitive applications in public clouds.

2. We formally define the two node deployment problems solved by ClouDiA and prove the hardness of these problems. The two optimization objectives used by ClouDiA—largest latency and longest critical path—model a large class of current distributed cloud applications that are sensitive to latency (Sect. 3).
3. We present an optimization framework that can solve these two classes of problems, and explore multiple algorithmic approaches. In addition to lightweight greedy and randomization techniques, we present different solvers for the two problems based on mixed-integer and constraint programming. We discuss optimizations and heuristics that allow us to obtain high-quality deployment plans over the scale of hundreds of instances (Sect. 4).
4. We discuss methods to obtain accurate latency measurements (Sect. 5) and evaluate our optimization framework with both synthetic and real distributed applications in Amazon EC2. We observe 15–55 % reduction in time-to-solution or response times. These benefits come exclusively from optimized deployment plans and require no changes to the specific application (Sect. 6).

This paper extends and subsumes its earlier conference version [71]. The substantial additional contribution of this paper is: (a) the exploration of lightweight algorithmic approaches to solve the node deployment problem under the two optimization objectives used by ClouDiA (Sects. 4.3 and 4.5 as well as experimental results in Sect. 6.5). This paper makes the following further additional contributions: (b) the exploration of additional metrics to model communication cost other than mean latency, namely mean latency plus standard deviation, and latency at the 99th percentile (Sect. 3.2 and experimental results in Sect. 6.4); (c) a discussion of overlapped execution of ClouDiA with target applications (Sect. 2.2); (d) the confirmation of the same effects of latency heterogeneity and mean latency stability in public cloud providers other than Amazon Web Services, namely Google Compute Engine and Rackspace Cloud Server (“Appendix 3”).

We discuss related work in Sect. 7 and then conclude.

2 Tuning for latency-sensitive applications in the cloud

To give a high-level intuition for our approach, we first describe the classes of applications we target in Sect. 2.1. We then describe the architecture that ClouDiA uses to suggest deployments for these applications in Sect. 2.2.

2.1 Latency-sensitive applications

We can classify latency-sensitive applications in the cloud into two broad classes: high-performance computing applications, for which the main performance goal is *time-to-*

solution, and service-oriented applications, for which the main performance goal is *response time for service calls*.

2.1.1 Goal: time-to-solution

A number of HPC applications simulate natural processes via long-running, distributed computations. For example, consider the simulation of collective animal movement published by Couzin et al. in Nature [20]. In this simulation, a group of animals, such as a fish school, moves together in a two-dimensional space. Animals maintain group cohesion by observing each other. In addition, a few animals try to influence the direction of movement of the whole group, e.g., because they have seen a predator or a food source. This simulation can be partitioned among multiple compute nodes through a spatial partitioning scheme [62]. At every time step of the simulation, neighboring nodes exchange messages before proceeding to the next time step. As the end of a time step is a logical barrier, worst-link latency essentially determines communication cost [1, 12, 35, 72]. Similar communication patterns are common in multiple linear algebra computations [22]. Another example of an HPC application where time-to-solution is critical is dynamic traffic assignment [64]. Here traffic patterns are extrapolated for a given time period, say 15 min, based on traffic data collected for the previous period. Simulation must be faster than real time so that simulation results can generate decisions that will improve traffic conditions for the next time period. Again, the simulation is distributed over multiple nodes, and computation is assigned based on a graph partitioning of the traffic network [64]. In all of these HPC applications, time-to-solution is dramatically affected by the latency of the worst link.

2.1.2 Goal: service response time

Web services and portals, as well as search engines, are prime cloud applications [6, 28, 60]. For example, consider a web portal, such as Yahoo! [52] or Amazon [60]. The rendering of a web page in these portals is the result of tens, or hundreds, of web service calls [44]. While different portions of the web page can be constructed independently, there is still a critical path of service calls that determines the server-side communication time to respond to a client request. Latencies in the critical path add-up and can negatively affect end user response time.

2.2 Architecture of ClouDiA

Figure 3 depicts the architecture of ClouDiA. The dashed line indicates the boundary between ClouDiA and public cloud tenants. The tuning methodology followed by ClouDiA comprises the following steps:

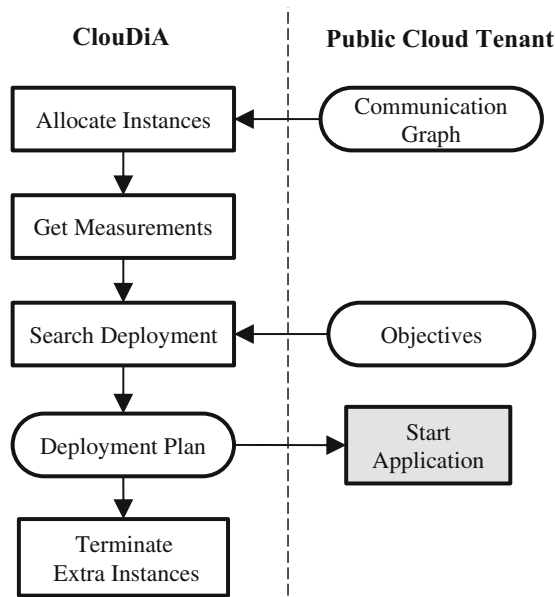


Fig. 3 Architecture of ClouDiA

1. **Allocate instances** A tenant specifies the communication graph for the application, along with a maximum number of instances at least as great as the required number of application nodes. ClouDiA then automatically allocates cloud instances to run the application. Depending on the specified maximum number of instances, ClouDiA will over-allocate instances to increase the chances of finding a good deployment.
2. **Get measurements** The pairwise latencies between instances can only be observed after instances are allocated. ClouDiA performs efficient network measurements to obtain these latencies, as described in Sect. 5. The main challenge is reliably estimating the mean latencies quickly, given that time spent in measurement is not available to the application.
3. **Search deployment** Using the measurement results, together with the optimization objective specified by the tenant, ClouDiA searches for a “good” deployment plan: one which avoids “bad” communication links. We formalize this notion and pose the node deployment problem in Sect. 3. We then formulate two variants of the problem that model our two classes of latency-sensitive applications. We prove the hardness of these optimization problems in “Appendix 1.” Given the hardness of these problems, traditional methods cannot scale to realistic sizes. We propose techniques that significantly speed up the search in Sect. 4.
4. **Terminate extra instances** Finally, ClouDiA terminates any over-allocated instances and the tenant can start the application with an optimized node deployment plan.

2.2.1 Adapting to changing network conditions

The architecture outlined above assumes that the application will run under relatively stable network conditions. We believe this assumption is justified: The target applications outlined in Sect. 2.1 have long execution time once deployed, and our experiments in Fig. 2 show stable pairwise latencies in EC2. In the future, more dynamic cloud network infrastructure may become the norm. In this case, the optimal deployment plan could change over time, forcing us to consider dynamic re-deployment. We envision that re-deployment can be achieved via iterations of the architecture above: getting new measurements, searching for a new optimal plan, and re-deploying the application.

Two interesting issues arise with iterative re-deployment. First, we need to consider whether information from previous runs could be exploited by a new deployment. Unfortunately, previous runs provide no information about network resources that were *not* used by the application. In addition, as re-deployment is triggered by changes in network conditions, it is unlikely that network conditions of previous runs will be predictive of conditions of future runs. Second, re-deployment should not interrupt the running application, especially in the case of web services and portals. Unfortunately, current public clouds do not support VM live migration [2,49,66]. Without live migration, complicated state migration logic would have to be added to individual cloud applications.

2.2.2 Overlapping ClouDiA with application execution

We envision one further improvement to our architecture that will become possible if support for VM live migration or state migration logic becomes pervasive: Instead of wasting idle compute cycles while ClouDiA performs network measurements and searches for a deployment plan, we could instead begin execution of the application over the initially allocated instances, in parallel with ClouDiA. Clearly, this strategy may lead to interference between ClouDiA’s measurements and the normal execution of the application, which would have to be carefully controlled. In addition, this strategy would only payoff if the state migration cost necessary to re-deploy the application under the plan found by ClouDiA would be small enough compared to simply running ClouDiA as outlined in Fig. 3.

3 The node deployment problem

In this section, we present the optimization problems addressed by ClouDiA. We begin by discussing how we model network cost (Sects. 3.1 and 3.2). We then formalize two versions of the node deployment problem (Sect. 3.3).

3.1 Cost functions

When a set of instances is allocated in a public cloud, in principle any instance in the set can communicate with any other instance, possibly at differing costs. Differences arise because of the underlying physical network resources that implement data transfers between instances. We call the collection of network resources that interconnect a pair of instances the *communication link* between them. We formalize communication cost as follows.

Definition 1 (*Communication Cost*) Given a set of instances S , we define $\mathcal{C}_{\mathcal{L}} : S \times S \rightarrow \mathbb{R}$ as the communication cost function. For a pair of instances $i, j \in S$, $\mathcal{C}_{\mathcal{L}}(i, j)$ gives the communication cost of the link from instance i to j .

$\mathcal{C}_{\mathcal{L}}$ can be defined based on different criteria, e.g., latency, bandwidth, or loss rate. To reflect true network properties, we assume costs of links can be asymmetric and the triangle inequality does not necessarily hold. In this paper, given the applications we are targeting, we focus solely on using *network latency* as a specific instance for $\mathcal{C}_{\mathcal{L}}$. Extending to other network cost measurements is future work.

Our definition of communication cost treats communication links as essentially independent. This modeling decision ignores the underlying implementation of communication links in the datacenter. In practice, however, current clouds tend to organize their network topology in a tree-like structure [11]. A natural question is whether we could provide more structure to the communication cost function by reverse engineering this underlying network topology. Approaches such as Sequoia [53] deduce underlying network topology by mapping application nodes onto leaves of a virtual tree. Unfortunately, even though these inference approaches work well for Internet-level topologies, state-of-the-art methods cannot infer public cloud environments accurately [10, 16].

Even if we could obtain accurate topology information, it would be non-trivial to make use of it. First, there is no guarantee that the nodes of a given application can all be allocated to nearby physical elements (e.g., in the same rack), so we need an approach that fundamentally tackles differences in communication costs. Second, datacenter topologies are themselves evolving, with recent proposals for new high-performance topologies [43]. Optimizations developed for a specific tree-like topology may no longer be generally applicable as new topologies are deployed. Our general formulation of communication cost makes our techniques applicable to multiple different choices of datacenter topologies. Nevertheless, as we will see in Sect. 6, we can support a general cost formulation, but optimize deployment search when there are uniformities in the cost, e.g., clusters of links with similar cost values.

To estimate the communication cost function $\mathcal{C}_{\mathcal{L}}$ for the set of allocated instances, ClouDiA runs an efficient network

measurement tool (Sect. 5). Note that the exact communication cost is application dependent. Applications communicating messages of different sizes can be affected differently by the network heterogeneity. However, for *latency-sensitive applications*, we expect that application-independent network latency measurements can be used as a good performance indicator, although they might not precisely match the exact communication cost in application execution. This is because our notion of cost need only discriminate between “good” and “bad” communication links, rather than accurately predict actual application runtime performance.

3.2 Metrics for communication cost

Even if we focus our attention solely on network latency, there are still multiple ways to measure and characterize such latency. The most natural metric, shown in Figs. 1 and 2, is mean latency, which captures the average latency behavior of a link. However, some applications are particularly sensitive to latency jitter, not only to heterogeneity in mean latency [72]. For these applications, an alternate metric which combines the mean latency with the standard deviation on latency measurements may be the most appropriate. Finally, demanding applications may seek latency guarantees at a high percentile of the latency distribution.

While all the above metrics provide genuine characterizations of different aspects of network latency, two additional considerations must be taken into account before adopting a latency metric. First, since in our framework communication cost is used merely to differentiate “good” from “bad” links, correlated metrics behave effectively in the same way. So a non-straight forward metric other than mean latency only makes sense if it is not significantly correlated with the mean. Second, any candidate latency metric should guide the search process carried out by ClouDiA such that lower-cost deployments represent deployments with lower actual time-to-solution or response time. Since most latency-sensitive applications tend to be significantly affected by mean latency, it is not clear whether other candidate latency metrics will lead to deployments with better application performance. We explore these considerations experimentally in Sect. 6.

3.3 Problem formulation

To deploy applications in public clouds, a mapping between logical application nodes and cloud instances needs to be determined. We call this mapping a *deployment plan*.

Definition 2 (*Deployment Plan*) Let S be the set of instances. Given a set N of application nodes, a deployment plan $\mathcal{D} : N \rightarrow S$ is a mapping of each application node $n \in N$ to an instance $s \in S$.

In this paper, we require that \mathcal{D} be an injection; that is, each instance $s \in S$ can have at most one application node $n \in N$ mapped to it. There are two implications of this definition. On the one hand, we do not collocate application nodes on the same instance. In some cases, it might be beneficial to collocate services, yet we argue these services should be merged into application nodes before determining the deployment plan. On the other hand, it is possible that some instances will have no application nodes mapped to them. This gives us flexibility to over-allocate instances at first and then shut-down those instances with high communication cost.

In today's public clouds, tenants typically determine a deployment plan by either a default or a random mapping. CloudDiA takes a more sophisticated approach. The deployment plan is generated by solving an optimization problem and searching through a space of possible deployment plans. This search procedure takes as input the communication cost function $\mathcal{C}_{\mathcal{L}}$, obtained by our measurement tool, as well as a *communication graph* and a *deployment cost function*, both specified by the cloud tenant.

Definition 3 (Communication Graph) Given a set N of application nodes, the communication graph $G = (V, E)$ is a directed graph where $V = N$ and $E = \{(i, j) | i, j \in N \wedge \text{talks}(i, j)\}$.

The *talks* relation above models the application-specific communication patterns. When defining the communication graph through the *talks* relation, the cloud tenant should only include communication links that have impact on the performance of the application. For example, those links first used for bootstrapping and rarely used afterwards should not be included in the communication graph. An alternative formulation for the communication graph would be to add weights to edges, extending the semantics of *talks*. We leave this to future work.

We find that although such a communication graph is typically not hard to extract from the application, it might be a tedious task for a cloud tenant to generate an input file with $O(|N|^2)$ links. CloudDiA therefore provides *communication graph templates* for certain common graph structures such as meshes or bipartite graphs to minimize human involvement.

In addition to the communication graph, a deployment cost function needs to be specified by the cloud tenant. At a high level, a deployment cost function evaluates the cost of the deployment plan by observing the structure of the communication graph and the communication cost for links in the given deployment. The optimization goal for CloudDiA is to generate a deployment plan that minimizes this cost. The formal definition of deployment cost is as follows:

Definition 4 (Deployment Cost) Given a deployment plan \mathcal{D} , a communication graph G , and a communication cost function $\mathcal{C}_{\mathcal{L}}$, we define $\mathcal{C}_{\mathcal{D}}(\mathcal{D}, G, \mathcal{C}_{\mathcal{L}}) \in \mathbb{R}$ as the deployment

cost of \mathcal{D} . $\mathcal{C}_{\mathcal{D}}$ must be monotonic on link cost and invariant under exchanging nodes that are indistinguishable using link costs.

Now, we can formally define the node deployment problem:

Definition 5 (Node Deployment Problem) Given a deployment cost function $\mathcal{C}_{\mathcal{D}}$, a communication graph G , and a communication cost function $\mathcal{C}_{\mathcal{L}}$, the node deployment problem is to find the optimal deployment

$$\mathcal{D}_{OPT} = \arg \min_{\mathcal{D}} \mathcal{C}_{\mathcal{D}}(\mathcal{D}, G, \mathcal{C}_{\mathcal{L}}).$$

In the remainder, we focus on two classes of deployment cost functions, which capture the essential aspects of the communication cost of latency-sensitive applications running in public clouds.

In high-performance computing (HPC) applications, such as simulations, matrix computations, and graph processing [72], application nodes typically synchronize periodically using either global or local communication barriers. The completion of these barriers depends on the communication link which experiences the longest delay. Motivated by such applications, we define our first class of deployment cost function to return the highest link cost.

Class 1 (Deployment Cost: Longest Link) Given a deployment plan \mathcal{D} , a communication graph $G = (V, E)$ and a communication cost function $\mathcal{C}_{\mathcal{L}}$, the longest link deployment cost $\mathcal{C}_{\mathcal{D}}^{LL}(\mathcal{D}, G, \mathcal{C}_{\mathcal{L}}) = \max_{(i, j) \in E} \mathcal{C}_{\mathcal{L}}(\mathcal{D}(i), \mathcal{D}(j))$.

Another class of latency-sensitive applications is exemplified by search engines [6, 8] as well as web services and portals [28, 44, 52, 60]. Services provided by these applications typically organize application nodes into trees or directed acyclic graphs, and the overall latency of these services is determined by the communication *path* which takes the longest time [58]. We therefore define our second class of deployment cost function to return the highest path cost.

Class 2 (Deployment Cost: Longest Path) Given a deployment plan \mathcal{D} , an acyclic communication graph G and a communication cost function $\mathcal{C}_{\mathcal{L}}$, the longest path deployment cost $\mathcal{C}_{\mathcal{D}}^{LP}(\mathcal{D}, G, \mathcal{C}_{\mathcal{L}}) = \max_{\text{path } P \subseteq G} (\sum_{(i, j) \in P} \mathcal{C}_{\mathcal{L}}(\mathcal{D}(i), \mathcal{D}(j)))$.

Note that the above definition assumes that the application is sending a sequence of causally related messages along the edges of a path, and summation is used to aggregate the communication cost of the links in the path.

Although we believe these two classes cover a wide spectrum of latency-sensitive cloud applications, there are still important applications which do not fall exactly into either of them. For example, consider a key-value store with latency requirements on average response time or the 99.9th percentile of the response time distribution [21]. This application does not exactly match either of the two deployment cost

functions above, since average response time may be influenced by multiple links in different paths. We discuss the applicability of our deployment cost functions to a key-value store workload further in Sect. 6.1. We then proceed to show experimentally in Sect. 6.4 that even though longest link is not a perfect match for such a workload, use of this deployment cost function still yields a 15–31 % improvement in average response time for a key-value store workload. Given the possibility of utilizing the above cost functions even if there is no exact match, CloudDiA is able to automatically improve response times of an even wider range of applications.

We prove that the node deployment problem with longest path deployment cost is NP-hard. With longest link deployment cost, it is also NP-hard and cannot be efficiently approximated unless $P = NP$ (“Appendix 1”). Given the hardness results, our solution approach consists of mixed-integer programming (MIP), constraint programming (CP) formulations, and lightweight algorithmic approaches (Sect. 4). We show experimentally that our approach brings significant performance improvement to real applications. In addition, from the insight gained from the theorems above, we also show that properly rounding communication costs to *cost clusters* can be heuristically used to further boost solver performance (Sect. 6.3).

4 Search techniques

In this section, we propose two encodings to solve the Longest Link Node Deployment Problem (LLNDP) using MIP and CP solvers (Sects. 4.1 and 4.2), as well as one formulation for the Longest Path Node Deployment Problem (LPNDP) using a MIP solver (Sect. 4.4). In addition to solver-based solutions, we also explore alternative lightweight algorithmic approaches to both problems (Sects. 4.3 and 4.5).

4.1 Mixed-integer program for LLNDP

Given a communication graph $G = (V, E)$ and a communication cost function $\mathcal{C}_{\mathcal{L}}$ defined over any pair of instances in S , the *Longest Link Node Deployment Problem* can be formulated as the following mixed-integer program (MIP):

$$(MIP) \text{ minimize } c$$

$$\text{s.t. } \sum_{i \in V} x_{ij} = 1 \quad \forall j \in S \quad (1)$$

$$\sum_{j \in S} x_{ij} = 1 \quad \forall i \in V \quad (2)$$

$$c \geq \mathcal{C}_{\mathcal{L}}(j, j')(x_{ij} + x_{i'j'} - 1)$$

$$\forall (i, i') \in E, \forall j, j' \in S$$

$$x_{ij} \in \{0, 1\} \quad \forall i \in V, j \in S$$

$$c \geq 0 \quad (3)$$

In this encoding, the boolean variable x_{ij} indicates whether the application node $i \in V$ is deployed on instance $j \in S$. The constraints (1) and (2) ensure that the variables x_{ij} represent a one-to-one mapping between the set V and S . Note that the set V might need to be augmented with dummy application nodes so that we have $|V| = |S|$. Also, the constraints (3) require that the value of c is at least equal to $\mathcal{C}_{\mathcal{L}}(j, j')$ any time there is a pair (i, i') of communicating application nodes and that i and i' are deployed on j and j' , respectively. Finally, the minimization in the objective function will make one of the constraints (3) tight, thus leading to the desired longest link value.

4.2 Constraint programming for LLNDP

Whereas the previous formulation directly follows from the problem definition, our second approach exploits the relation between this problem and the subgraph isomorphism problem, as well as the clusters of communication cost values. The algorithm proceeds as follows. Given a goal c , we search for a deployment that avoids communication costs greater than c . Such a deployment exists if and only the graph $G_c = (S, E_c)$ where $E_c = \{(i, j) : \mathcal{C}_{\mathcal{L}}(i, j) \leq c\}$ contains a subgraph isomorphic to the communication graph $G = (V, E)$. Therefore, by finding such a subgraph, we obtain a deployment whose deployment cost c' is such that $c' \leq c$. Assume c'' is the largest communication cost strictly lower than c' . Thus, any improving deployment must have a cost of c'' or lower, and the communication graph $G = (V, E)$ must be isomorphic to a subgraph of $G_{c''} = (S, E_{c''})$ where $E_{c''} = \{(i, j) : \mathcal{C}_{\mathcal{L}}(i, j) \leq c''\}$. We proceed iteratively until no deployment is found. Note that the number of iterations is bounded by the number of distinct cost values. Therefore, clustering similar values to reduce the number of distinct cost values would improve the computation time by lowering the number of iterations, although it approximates the actual value of the objective function. We investigate the impact of *cost clusters* in Sect. 6. For a given objective value c , the encoding of the problem might be expressed as the following constraint programming (CP) formulation:

$$(CP) \text{ alldifferent}((u_i)_{1 \leq i \leq |V|})$$

$$(u_i, u_{i'}) \neq (j, j') \quad \forall (i, i') \in E, \forall j, j' \in S :$$

$$\mathcal{C}_{\mathcal{L}}(j, j') > c$$

$$u_i \in \{1, \dots, |S|\} \quad \forall 1 \leq i \leq |V|$$

This encoding is substantially more compact than the MIP formulation, as the binary variables are replaced by integer variables, and the mapping are efficiently captured within the `alldifferent` constraint. In addition, at the root of the search tree, we perform an extra filtering of the domains of

Algorithm 1 G1

Require: Instances S
Require: Communication Graph $G = (V, E)$

- 1: Find $(u_0, v_0) \in S \times S$ of lowest cost
- 2: Find an arbitrary edge $(x, y) \in E$
- 3: $\mathcal{D}(x) = u_0, \mathcal{D}(y) = v_0$
- 4: **for** $i = 1$ to $|V| - 2$ **do**
- 5: $c_{\min} = \infty$
- 6: **for** $(u, v) \in S \times S$ **do**
- 7: **if** $\mathcal{D}^{-1}(v)$ is undefined and
 $\mathcal{D}^{-1}(u)$ has unmatched neighbors **then**
- 8: **if** $\mathcal{C}_{\mathcal{L}}(u, v) < c_{\min}$ **then**
- 9: $c_{\min} = \mathcal{C}_{\mathcal{L}}(u, v)$
- 10: $u_{\min} = u, v_{\min} = v$
- 11: **end if**
- 12: **end if**
- 13: **end for**
- 14: $w =$ one of $\mathcal{D}^{-1}(u_{\min})$'s unmatched neighbors
- 15: $\mathcal{D}(w) = v_{\min}$
- 16: **end for**

the x_{ij} variables that is based on compatibility between application nodes and instances. Indeed, as the objective value c decreases, the graph G_c becomes sparse, and some application nodes can no longer be mapped to some of the instance nodes. For example, a node in G needs to be mapped to a node in G_c of equal or higher degree. Similar to [70], we define a labeling based on in- and out-degree, as well as information about the labels of neighboring nodes. This labeling establishes a partial order on the nodes and expresses compatibility between them. This compatibility relationship is used to remove from the domain of an application node every instance that would not be compatible. For more details on this labeling, please refer to [70].

4.3 Lightweight approaches for LLNDP

Randomization and greedy approaches can also be applied to the LLNDP. We explore each of these approaches in turn.

4.3.1 Randomization

The easiest approach to finding a (suboptimal) solution for LLNDP is to generate a number of deployments randomly and select the one with the lowest deployment cost. Compared with CP or MIP solutions, generating deployments randomly explores the search space in a less intelligent way. However, since generating deployments is computationally cheaper and easier to parallelize, it is possible to explore a larger portion of the search space given the same amount of time.

4.3.2 Greedy algorithms

Greedy algorithms can also be used as lightweight approaches to quickly find a (suboptimal) solution for LLNDP. We present two greedy approaches:

Algorithm 2 G2

Require: Instances S
Require: Communication Graph $G = (V, E)$

- 1: Find $(u_0, v_0) \in S \times S$ of lowest cost
- 2: Find an arbitrary edge $(x, y) \in E$
- 3: $\mathcal{D}(x) = u_0, \mathcal{D}(y) = v_0$
- 4: **for** $i = 1$ to $|V| - 2$ **do**
- 5: $c_{\min} = \infty$
- 6: **for** $(u, v) \in S \times S$ **do**
- 7: **for** w where $(\mathcal{D}^{-1}(u), w) \in E$ **do**
- 8: $c_{uv} = \mathcal{C}_{\mathcal{L}}(u, v)$
- 9: **for** $(w, x) \in E$ **do**
- 10: **if** $\mathcal{D}(x)$ is defined and $\mathcal{C}_{\mathcal{L}}(v, \mathcal{D}(x)) > c_{uv}$ **then**
- 11: $c_{uv} = \mathcal{C}_{\mathcal{L}}(v, \mathcal{D}(x))$
- 12: **end if**
- 13: **end for**
- 14: **if** $c_{uv} < c_{\min}$ **then**
- 15: $c_{\min} = c_{uv}$
- 16: $v_{\min} = v, w_{\min} = w$
- 17: **end if**
- 18: **end for**
- 19: **end for**
- 20: $\mathcal{D}(w_{\min}) = v_{\min}$
- 21: **end for**

(G1) Recall that the deployment plan \mathcal{D} is a mapping from application nodes to instances. Let \mathcal{D}^{-1} be the inverse function of \mathcal{D} , mapping each instance $s \in S$ to an application node $n \in V$. The first greedy approach, shown in Algorithm 1, works as follows:

1. Find a link $(u_0, v_0) \in S \times S$ of lowest cost, and for an arbitrary edge $(x, y) \in E$, let $\mathcal{D}(x) = u_0, \mathcal{D}(y) = v_0$ (Lines 1–3);
2. Find a link $(u, v) \in S \times S$ of lowest cost s.t. instance u is mapped to a node in the current partial deployment that still has unmatched neighbors, and instance v is not mapped in the current deployment (Lines 5–13);
3. Add the instance v to the partial deployment by letting $\mathcal{D}(w) = v$, where $(\mathcal{D}^{-1}(u), w)$ is one of the unmapped edges in E (Lines 14 and 15);
4. Repeat Steps 2 and 3 until all nodes are included (Lines 4–16).

This greedy approach is simple and intuitive, but it has one potential drawback: Although the links explicitly picked by the algorithm typically have low cost, the implicit links introduced while selecting a partial solution following the lowest cost-edge criterion can have substantial cost. This is because the mapping of a node to an instance v implies that other nodes already in the deployment are then connected to v by the corresponding underlying links. We address this issue in the refined greedy approach below.

(G2) In order to avoid selecting high-cost links implicitly when mapping a low-cost link, we revise the lowest cost-edge criterion in Step 2 above as shown in Algorithm 2. Instead of

costing a particular $(u, v) \in S \times S$ simply by the cost of the corresponding link, we take the highest cost among the cost of (u, v) and of all links between $D(w)$ and v assuming node w is added to the current partial deployment (Lines 7–18). Intuitively, we consider not only the explicit cost of a given link that is a candidate for addition to the deployment, but also the costs of all other links which would be implicitly added to the deployment by this candidate mapping. By selecting the candidate with the minimum cost among both explicit and implicit link additions, this greedy variant locally minimizes the longest link objective at each decision point.

4.4 Mixed-integer programming for LPNDP

As previously mentioned, the node deployment problem is intrinsically related to the subgraph isomorphism problem (SIP). In addition, the longest link objective function allows us to directly prune the graph that must contain the communication graph G and therefore can be encoded as a series of subgraph isomorphism satisfaction problems. This plays a key role in the success of the CP formulation. On the contrary, the objective function of the *Longest Path Node Deployment Problem* (LPNDP) interferes with the structure of the SIP problem and rules out sub-optimal solutions only when most of the application nodes have been assigned to instances. As a result, this optimization function barely guides the systematic search and makes it less suitable for a CP approach. Consequently, we only provide a MIP formulation for the LPNDP.

Given a communication graph $G = (V, E)$ and a communication cost function $C_{\mathcal{L}}$ defined over any pair of instances in S , the *Longest Path Node Deployment Problem* (LPNDP) can be formulated as the following mixed-integer program (MIP):

$$\begin{aligned}
 & \text{(MIP) minimize } t \\
 & \text{s.t. } \sum_{i \in V} x_{ij} = 1 \quad \forall j \in S \\
 & \quad \sum_{j \in S} x_{ij} = 1 \quad \forall i \in V \\
 & \quad c_{ii'} \geq C_{\mathcal{L}}(j, j')(x_{ij} + x_{i'j'} - 1) \quad \forall (i, i') \in E, \forall j, j' \in S \\
 & \quad t \geq t_i, t_i \geq 0 \quad \forall i \in V \\
 & \quad t_{i'} \geq t_i + c_{ii'} \quad \forall (i, i') \in E \\
 & \quad x_{ij} \in \{0, 1\} \quad \forall i \in V, j \in S \\
 & \quad c_{ii'} \geq 0 \quad \forall (i, i') \in E \\
 & \quad t \geq 0
 \end{aligned}$$

As in the previous MIP encoding, the boolean variable x_{ij} indicates whether the application node i is deployed on instance j in a one-to-one mapping. In addition, the variable $c_{ii'}$ captures the communication cost from application node i to node i' that would result from the deployment specified

by the x_{ij} variables. The variable t_i represents the longest directed path in the communication graph G that reaches the application node i . Finally, the variable t appears in the objective function and corresponds to the maximum among the t_i variables.

4.5 Lightweight approaches for LPNDP

As with LLNDP, we explore both randomization and greedy approaches.

4.5.1 Randomization

Similarly to LLNDP, we can find a (suboptimal) solution for LPNDP by generating a number of random deployments in parallel and selecting one with the lowest deployment cost.

4.5.2 Greedy heuristic approach

Since the communication graph for LPNDP can be any directed acyclic graph containing paths of different lengths, the effect of adding a single node to a given partial deployment cannot be easily estimated. Therefore, the greedy algorithms described in Sect. 4.3 cannot be directly extended to LPNDP. However, given a LPNDP with communication graph G , we can still solve LLNDP with G greedily and use the resulting mapping as a heuristic solution for LPNDP. We experimentally study the effectiveness of these lightweight approaches in Sect. 6.5.

5 Measuring network distance

Making wise deployment decisions to optimize performance for latency-sensitive applications requires knowledge of pairwise communication cost. A natural way to characterize the communication cost is to directly measure round-trip latencies for all instance pairs. To ensure such latencies are a good estimate of communication cost during application execution, two aspects need to be handled. First, the size of the messages being exchanged during application execution is usually nonzero. Therefore, rather than measuring pure round-trip latencies with no data content included, we measure TCP round-trip time of small messages, where message size depends on the actual application workload. Second, during the application execution, multiple messages are typically being sent and received at the same time. Such temporal correlation affects network performance, especially end-to-end latency. But the exact interference patterns heavily depend on low-level implementation details of applications, and it is impractical to require such detailed information from the tenants. Instead, we focus on estimating the quality of links

without interference, as this already gives us guidance on which links are certain to negatively affect actual executions.

Experimental studies have demonstrated that clouds suffer from high latency jitter [56,61,72]. However, we have observed experimentally that it is possible to obtain stable measurements of mean latency, if sufficient repetitions are carried out (see results in Fig. 2). Unfortunately, to estimate mean latency accurately and properly capture latency heterogeneity among links, multiple round-trip latency measurements have to be obtained for each pair of instances. Since the number of instance pairs is quadratic in the number of instances, such measurement takes substantial time. On the one hand, we want to run mean-latency measurements as fast as possible to minimize the overhead of using ClouDiA. On the other hand, we need to avoid introducing uncontrolled measurement artifacts that may affect the quality of our results. We propose three possible approaches for organizing pairwise mean latency measurements in the following.

1. **Token Passing.** In this first approach, a unique token is passed between instances. When an instance i receives this token, it selects another instance j and sends out a probe message of given size. Once the entire probe message has been received by j , it replies to i with a message of the same size. Upon receiving the entire reply message, i records the round-trip time and passes the token on to another instance chosen at random or using a predefined order. By having such a unique token, we ensure that only one message is being transferred at any given time, including the message for token passing itself. We repeat this token passing process a sufficiently large number of times, so multiple round-trip measurements can be collected for each link. We then aggregate these measurements into mean latencies per link.

This approach achieves the goal of obtaining pairwise mean-latency measurements without correlations between links. However, the lack of parallelism restricts its scalability.

2. **Uncoordinated.** To improve scalability, we would like to avoid excessive coordination among instances, so that they can execute measurements in parallel. We introduce parallelism by the following simple scheme: Each instance picks a destination at random and sends out a probe message. Meanwhile, all instances monitor incoming messages and send reply messages once an entire probe message has been received. After one such round-trip measurement, each instance picks another probe destination and starts over. The process is repeated until we have collected enough round-trip measurements for every link. We then aggregate these measurements into mean latencies per link.

Given n instances, this approach allows up to n messages to be in flight at any given time. Therefore, this approach

provides better scalability than the first approach. However, since probe destinations are chosen at random without coordination, it is possible that: 1) one instance needs to send out reply messages while it is sending out a probe message; or 2) multiple probe messages are sent to the same destination from different sources. Such cross-link correlations are undesirable for our measurement goal.

3. **Staged.** To prevent cross-link correlations while preserving scalability, coordination is required when choosing probe destinations. We add an extra *coordinator* instance and divide the entire measurement process into *stages*. To start a stage for n instances in parallel, the coordinator first picks $\lfloor \frac{n}{2} \rfloor$ pairs of instances $\{(i_1, j_1), (i_2, j_2), \dots, (i_{\lfloor \frac{n}{2} \rfloor}, j_{\lfloor \frac{n}{2} \rfloor})\}$ such that $\forall p, q \in \{1..n\}, i_p \neq j_q$ and $i_p \neq i_q$ if $p \neq q$. The coordinator then notifies each $i_p, p \in \{1, \dots, \lfloor \frac{n}{2} \rfloor\}$, of its corresponding j_p . After receiving a notification, i_p sends probe messages to j_p and measures round-trip latency as described above. Finally, i_p ends its stage by sending a notification back to the coordinator, and the coordinator waits for all pairs to finish before starting a new stage.

This approach allows up to $\frac{n}{2}$ messages between instances in flight at any time at the cost of having a central coordinator. We minimize the cost of per-stage coordination by consecutively measuring round-trip times between the same given pair of instances K_s times within the same stage, where K_s is a parameter. With this optimization, the staged approach can potentially provide scalability comparable to the *uncoordinated* approach. At the same time, by careful implementation, we can guarantee that each instance is always in one of the following three states: (1) sending to one other instance; (2) receiving from one other instance; or (3) idle in networking. This guarantee provides independence among pairwise link measurements similar to that achieved by *token passing*.

5.1 Approximations

Even the staged network latency benchmark above can take non-negligible time to generate mean latency estimates for a large number of instances. Given that our goal is simply to estimate link costs for our solvers, we have experimented with other simple network metrics, such as hop count and IP distance, as proxies for round-trip latency. Surprisingly, these metrics did *not* turn out to correlate well with round-trip latency. We provide details on these negative results in “Appendix 2.”

6 Experimental results

In this section, we present experiments demonstrating the effectiveness of ClouDiA. We begin with a description of

the several representative workloads used in our experiments (Sect. 6.1). We then present micro-benchmark results for the network measurement tools and the solver techniques of ClouDiA (Sects. 6.2 and 6.3). Next, we experimentally demonstrate the performance improvements achievable in public clouds by using ClouDiA as the deployment advisor (Sect. 6.4). Finally, we show the effectiveness of lightweight algorithmic approaches compared with solver-based solutions (Sect. 6.5).

6.1 Workloads

To benchmark the performance of ClouDiA, we implement three different workloads: a behavioral simulation workload, a query aggregation workload, and a key-value store workload. Each workload illustrates a different communication pattern.

6.1.1 Behavioral simulation workload

In behavioral simulations, collections of individuals interact with each other to form complex systems. Examples of behavioral simulations include large-scale traffic simulations and simulation of groups of animals. These simulations organize computation into ticks and achieve parallelism by partitioning the simulated space into regions. Each region is allocated to a processor and internode communication is organized as a 2D or 3D mesh. As synchronization among processors happens every tick, the progress of the entire simulation is limited by the pair of nodes that take longest to synchronize. Longest link is thus a natural fit to the deployment cost of such applications. We implement a workload similar to the fish simulation described by Couzin et al. [20]. The communication graph is a 2D mesh, and the message size per link is 1 KB for each tick. To focus on network effects, we hide CPU-intensive computation and study the time to complete 100K ticks over different deployments.

6.1.2 Synthetic aggregation query workload

In a search engine or distributed text database, queries are processed by individual nodes in parallel and the results are then aggregated [8]. To prevent the aggregation node from becoming a bottleneck, a multi-level aggregation tree can be used: Each node aggregates some results and forwards the partial aggregate to its parent in the tree for further aggregation. The response time of the query depends on the path from a leaf to the root that has highest total latency. Longest path is thus a natural fit for the deployment cost of such applications. We implement a two-level aggregation tree of a top-k query answering workload. The communication graph is a tree and the forwarding message size varies from the leaves to the root, with an average of 4 KB. We hide ranking score com-

putation and study the response time of aggregation query results over different deployments.

6.1.3 Key-value store workload

We also implement a distributed key-value store workload. The key-value store is queried by a set of front-end servers. Keys are randomly partitioned among the storage nodes, and each query touches a random subset of them. The communication graph therefore is a bipartite graph between front-end servers and storage machines. However, unlike the simulation workload, the average response time of a query is not simply governed by the slowest link. To see this, consider a deployment with mostly equal-cost links, but with a single slower link of cost c , and compare this to a similar deployment with two links of cost $c - \epsilon$. If longest link were used as the deployment cost function, the second deployment would be favored even though the first deployment actually has lower average response time. Indeed, neither longest link nor longest path is the precisely correct objective function for this workload. We evaluate the query response time over different deployments by using longest link, with a hope that it can still help avoid high-cost links.

6.2 Network micro-benchmarks

6.2.1 Setup

The network measurement tools of ClouDiA are implemented in C++ using TCP sockets (SOCK_STREAM). We set all sockets to non-blocking mode and use *select* to process concurrent flows (if any). We disable the Nagle Algorithm.

We ran experiments in the Amazon Elastic Compute Cloud (Amazon EC2). We used large instances (m1.large) in all experiments. Each large instance has 7.5 GB memory and four EC2 Compute Units. Each EC2 Compute Unit provides the equivalent CPU capacity of a 1.0–1.2 GHz 2007 Opteron or 2007 Xeon processor [2]. Unless otherwise stated, we use 100 large instances for network micro-benchmarks, all allocated by a single *ec2-run-instance* command, and set the round-trip message size to 1 KB. We show the following results from the same allocation so that they are comparable. Similar results are obtained in other allocations.

6.2.2 Round-trip latency measurement

We run latency measurements with each of the three approaches proposed in Sect. 5 and compare their accuracy in Fig. 4. To make sure *token passing* can observe each link a sufficiently large number of times, we use 50 large instances in this experiment. We consider the mean latencies for 50^2 instance pairs as a 50^2 -dimension vector of mean latencies, of which each dimension represents one link. In ClouDiA,

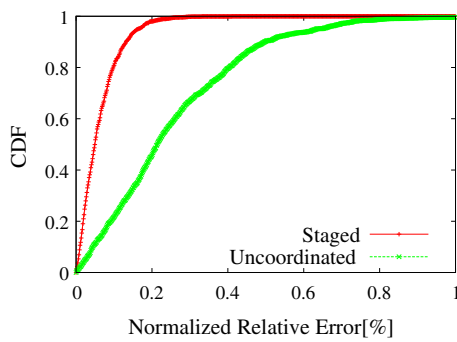


Fig. 4 Normalized relative error of network distance measurement methods against token passing. Staged shows higher accuracy than uncoordinated

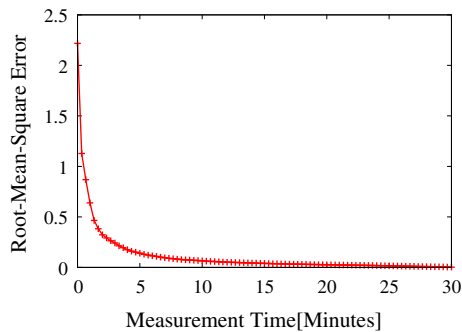


Fig. 5 Latency measurement convergence over time. Root-mean-square error drops quickly within the first 5 min

link latencies are only used to determine the relative order in choosing links or paths; thus, if one methodology overestimates or underestimates all link latencies by the same factor, its measurements result in the same deployment plan as the ideal plan generated from the accurate measurements. To avoid counting such overestimation or underestimation as error, measurements are first normalized to the unit vector. Then, *staged* and *uncoordinated* are compared with the baseline *token passing*. The CDF of the relative error of each dimension is shown in Fig. 4. Using *staged*, we find 90% of links have less than 10% relative error and the maximum error is less than 30%, whereas using *uncoordinated* we find 10% of links have more than 50% relative error. Therefore, as expected, *staged* exhibits higher measurement accuracy than *uncoordinated*.

Figure 5 shows convergence over time using the *staged* approach with 100 instances and $K_s = 10$. Again, the measurement result is considered as a latency vector. The result of the full 30 min observation is used as the *ground truth*. Each stage on average takes 2.75 ms. Therefore, we obtain about 3,004 measurements for each instance pair within 30 min. We then calculate the root-mean-square error of partial observations between 1 and 30 min compared with the ground truth. From Fig. 5, we observe the root-mean-square error drops quickly within the first 5 min and smooths out afterwards.

Therefore, we pick 5 min as the measurement time for all the following experiments with 100 instances. For experiments with $n \neq 100$ instances, since the *staged* approach tests $\frac{n}{2}$ pairs in parallel whereas there are $O(n^2)$ total pairs, measurement time needs to be adjusted linearly to $5 \cdot \frac{n}{100}$ minutes.

6.3 Solver micro-benchmarks

6.3.1 Setup

We solve the MIP formulation using the IBM ILOG CPLEX Optimizer, while every iteration of the CP formulation is performed using IBM ILOG CP Optimizer. Solvers are executed on a local machine with 12 GB memory and Intel Core i7-2600 CPU (four physical cores with hyper-threading). We enable parallel mode to allow both solvers to fully utilize all CPU resources. We use *k*-means to cluster link costs. Since the link costs are in one dimension, such *k*-means can be optimally solved in $O(kN)$ time using dynamic programming, where N is the number of distinct values for clustering and k is the number of *cost clusters*. After running *k*-means clustering, all costs are modified to the mean of the containing cluster and then passed to the solver. For comparison purposes, the same set of network latency measurements as in Sect. 6.2 is used for all solver micro-benchmarks. In each solver micro-benchmark, we set the number of application nodes to be 90% of the number of allocated instances. To find an initial solution to bootstrap the solver's search, we randomly generate 10 node deployment plans and pick the best one among those.

6.3.2 Longest link node deployment problem

We provide both a MIP formulation and a CP formulation for LLNDP. Since the parameter of cost clustering may have different impacts on the performance of the two solvers, we first analyze the effect of cost clustering. Figure 6 shows the convergence of the CP formulation for 100 instances with different numbers of clusters. The communication graph for Figs. 6, 7, and 8 is a 2D mesh from the simulation workload, and the deployment cost is the cost of the longest link. We tested all possible k values from 5 to the number of distinct values (rounded to nearest 0.01 ms), with an increment of 5. We present three representative configurations: $k = 5$, $k = 20$, and *no clustering*. As we decrease the number of clusters, the CP approach converges faster. Indeed, with no clustering, the best solution is found after 16 min, whereas it takes 2 min and 30 s for the CP approach to converge with $k = 20$ and $k = 5$, respectively. This is mainly due to the fact that fewer iterations are needed to reach the optimal value. Such a difference demonstrates the effectiveness of cost clustering in reducing the search time. On the other hand,

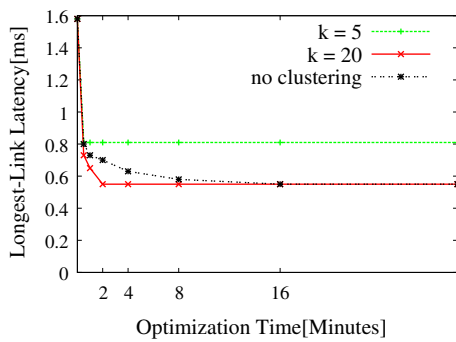


Fig. 6 Convergence for solving LLNDP using CP with different number of cost clusters. $k = 20$ converges faster than $k = 5$ and *no clustering*

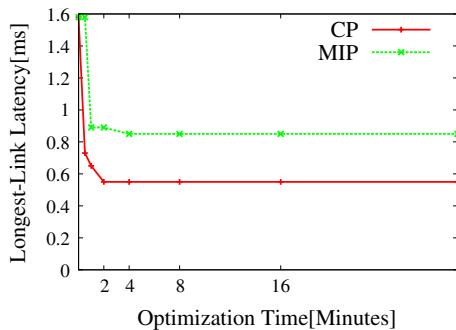


Fig. 7 Convergence for solving LLNDP using CP and MIP with 20 cost clusters. CP finds a significantly better solution

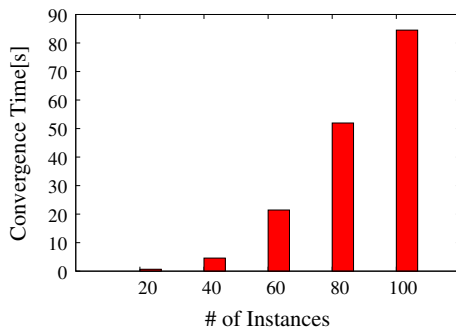


Fig. 8 Scalability for solving LLNDP using CP. Average convergence time increases acceptably with number of instances

the smaller the value of k is, the coarser the cost clusters are. As a result, the CP model cannot discriminate among deployments within the same cost cluster, and this might lead to sub-optimal solutions. As shown in Fig. 6, the solver cannot find a solution with a deployment cost smaller than 0.81 for $k = 5$, while both cases $k = 20$ and *no clustering* lead to a much better deployment cost of 0.55.

Figure 7 shows the comparison between the CP and the MIP formulations with $k = 20$. MIP performs poorly with the scale of 100 instances. Also, other clustering configurations do not improve the performance of MIP. One reason is that for LLNDP, the encoding of the MIP is much less compact than CP. Moreover, the MIP formulation suffers from a weak

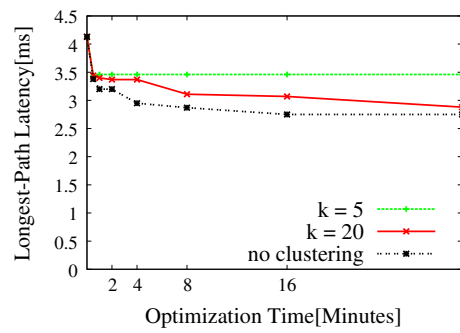


Fig. 9 Convergence for solving LPNDP using MIP. Clustering does not improve performance

linear relaxation, as x_{ij} and $x_{i'j'}$ should add-up to more than one for the relaxed constraint 3 to take effect.

Given the above analysis, we pick CP with $k = 20$ for the following scalability experiments as well as LLNDP in Sects. 6.4 and 6.5.

Measuring scalability of a solver such as ours is challenging, as problem size does not necessarily correlate with problem hardness. To observe scalability behavior with problem size, we generate multiple inputs for each size and measure the average convergence time of the solver over all inputs. The multiple inputs for each size are obtained by randomly choosing 50 subsets of instances out of our initial 100-instance allocation. The convergence time corresponds to the time the solver takes to not be able to improve upon the best found solution within 1 h of search. Figure 8 shows the scalability of the solver with the CP formulation. We observe that average convergence time increases acceptably with the problem size. At the same time, at every size, the solver is able to devise node deployment plans with similar average deployment cost improvement ratios.

6.3.3 Longest path node deployment problem

Figure 9 shows the convergence of the MIP formulation for 50 instances with different number of link cost clusters. The communication graph is an aggregation tree with depth less than or equal to 4. Similar to Fig. 6, the solver performs poorly under the configuration of $k = 5$. Interestingly, clustering costs does not improve performance for LPNDP. This is because the costs are aggregated using summation over the path for LPNDP, and therefore, the solver cannot take advantage of having fewer distinct values. We therefore use MIP with *no clustering* for LPNDP in Sects. 6.4 and 6.5.

6.4 ClouDiA effectiveness

6.4.1 Setup

We evaluate the overall ClouDiA system in EC2 with 100–150 instances over different allocations. Other settings are the

same as in Sect. 6.2 (network measurement) and Sect. 6.3 (solver). We use a 10% over-allocation ratio in all experiments except the last one (Fig. 12), in which we vary this ratio. The deployment decision made by ClouDiA is compared with the default deployment, which uses the instance ordering returned by the EC2 allocation command. Note that EC2 does not offer us any control on how to place instances, so the over-allocated instances we obtain are just the ones returned by the `ec2-run-instance` command.

6.4.2 Cost metrics

In Fig. 10, we study the correlation between three communication cost metrics under one representative allocation of 110 instances. Each point represents the link between a pair of nodes: The x axis shows its mean latency (mean) and the y axis shows its mean latency plus standard deviation (mean+SD) or 99th percentile latency (99%). While links with larger mean latencies tend to have larger mean+SD or 99% values, they are not perfectly correlated. This result motivates us to study the actual application performance under deployments generated by ClouDiA with different cost metrics. Figure 11 shows the relative improvement of using Mean + SD or 99% compared with using mean. Using 99th percentile latency reduces actual performance for all three applications, suggesting that the performance of these applications is not well-captured solely by this metric. While using Mean+SD improves performance for behavioral simulation and aggregation query workloads, it reduces performance for the key-value store workload. However, the observed differences in performance with respect to using mean are not dramatic. These results suggest that even with the existence of latency jitter in public clouds, mean latency is still a robust metric for these three applications, which are sensitive to both mean latency and latency jitter. However, it is possible for an application to be insensitive to mean latency while remaining sensitive to latency jitter (e.g., an application that can completely hide latency when there is no jitter). For such an application, latency metrics other than mean latency may work better.

6.4.3 Overall effectiveness

We show the overall percentage of improvement over five different allocations in EC2 in Fig. 12. The behavioral simulation and key-value store workloads use 100 application nodes, whereas the aggregation query workload uses 50 nodes. For the simulation workload, we report the reduction in time-to-solution. For the aggregation query and key-value store workloads, we report the reduction in response time. Each of these is averaged based on an at least 10 min of observation for both. We compare the performance of ClouDiA optimized deployment to the default deployment. ClouDiA

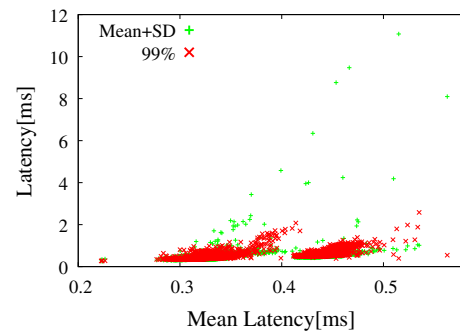


Fig. 10 Correlation between different cost metrics under one representative allocation of 110 instances. The cost metrics are not perfectly correlated with mean latency

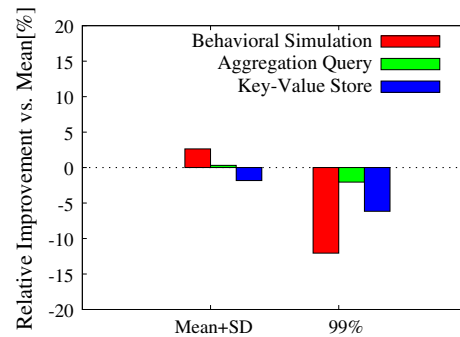


Fig. 11 Relative performance improvement of using other cost metrics compared with using mean latency. Mean latency is a robust metric for the applications considered

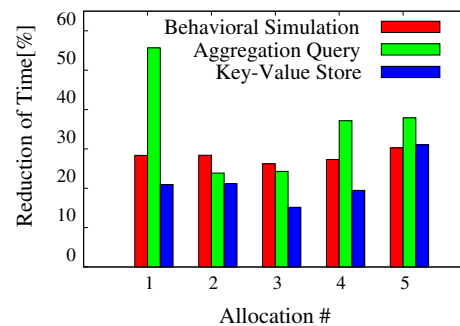


Fig. 12 Time reduction in EC2 over five different allocations. Reduction ratio is between 15 and 55%

achieves 15–55% reduction in time-to-solution or response time over 5 allocations for three workloads. The reduction ratio varies for different allocations. Among three workloads, we observe the largest reduction ratio on average in the aggregation query workload, while the key-value store workload gets less improvement than the others. There are two reasons for this effect. First, the communication graph of the aggregation query workload has the fewest edges, which increases the probability that the solver can find a high-quality deployment. Second, the longest link deployment cost function does not exactly match the mean response time measurement of the

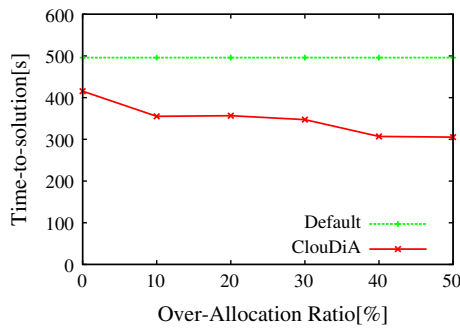


Fig. 13 Time-to-solution reduction for behavioral simulation with different over-allocation ratios. The largest improvement is obtained with the first 10 % of over-allocation

key-value store workload, and therefore, deployment decisions are made less accurately.

6.4.4 Effect of over-allocation

In Fig. 13, we study the benefit of over-allocating instances for increasing the probability of finding a good deployment plan. Note that although ClouDiA terminates extra instances once the deployment plan is determined, these instances will still be charged for at least 1 h usage due to the roundup pricing model used by major cloud service providers [2, 49, 66]. Therefore, a trade-off must be made between performance and initial allocation cost. In this experiment, we use an application workload similar to Fig. 12, but with 150 EC2 instances allocated at once by a single *ec2-run-instance command*. To study the case with over-allocation ratio x , we use the first $(1 + x) \cdot 100$ instances out of the 150 instances by the EC2 default ordering. Figure 13 shows the improvement in time-to-solution for the simulation workload. The default deployment always uses the first 100 instances, whereas ClouDiA searches deployment plans with the $100 \times$ extra instances. We report 38 % performance improvement with 50 % extra instances over-allocated. Without any over-allocation, 16 % improvement is already achieved by finding a good injection of application nodes to instances. Interestingly, with only 10 % instance over-allocation, 28 % improvement is achieved. Similar observations are found on other allocations as well.

6.5 Lightweight approaches

6.5.1 Setup

In Figs. 14 and 15, we compare the effectiveness of lightweight approaches against the CP and MIP formulations. Results are averaged over 20 different allocations of 50 instances with 10 % over-allocation. G1 is the simple greedy algorithm, which adds a node following a lowest cost-edge criterion at each step. G2 is the refined greedy algorithm,

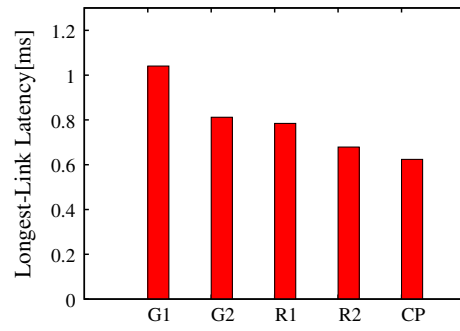


Fig. 14 Effectiveness of using lightweight approaches against CP for LLNDP. R2 finds a solution with longest link latency only 8.65 % higher than CP

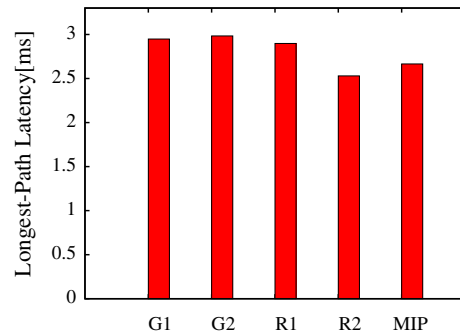


Fig. 15 Effectiveness of using lightweight approaches against MIP for LPNDP. R2 finds a solution with longest path latency 5.10 % lower than MIP

which iteratively adds a node such that partial deployment cost is minimal after addition. R1 is the lowest deployment cost obtained by generating 1,000 random deployments. R2 is the lowest deployment cost obtained by generating random deployments in parallel using the same amount of wall-clock time as well as the same hardware given to the CP or MIP solvers. The solver setup and hardware configuration are the same as in Sect. 6.3.

6.5.2 Longest link node deployment problem

In Fig. 14, both CP and R2 run for 2 min and all other methods finish in less than 1 s. G1 provides the worst solution overall, with a cost 66.7 % higher than CP. We examine the top-5 implicit links that are not explicitly chosen by G1, but are nevertheless included in the final solution. These links are on average 31.6 % more expensive than the worst link picked by CP. G2 improves G1 significantly by taking implicit links into consideration during each step of solution expansion. Interestingly, R1 is able to generate deployments with average cost 3.39 % lower than G2. R2 is able to generate deployments with cost only 8.65 % higher than the best deployment found by CP. Such results suggest that simply generating a large number of random deployments and picking the best one can provide reasonable effectiveness with minimal development overhead.

6.5.3 Longest path node deployment problem

In Fig. 15, both MIP and R2 run for 15 min and all other methods finish in less than 1 s. Although G1 and G2 are designed for LLNDP, they are still able to generate deployments with cost comparable to R1. Surprisingly, R2 is able to find deployments with cost on average 5.10 % lower than MIP. We conjecture that even though the MIP solver can exploit the search space in a much more intelligent way than R2, the distribution of good solutions in this particular problem makes such intelligent searching less important. Meanwhile, R2's simplicity and efficiency enable it to explore a larger portion of the search space than MIP within the same amount of time.

To verify the effectiveness of R2, we also ran an additional experiment where the total number of instances was decreased to 15. In this scenario, MIP was always able to find optimal solutions within 15 min over 20 different allocations. Meanwhile, R2 found suboptimal solutions for 40 % of the allocations given the same amount of time as MIP. So we argue that MIP is still a complete solution which guarantees optimality when the search finishes. R2, however, cannot provide any guarantee, even when the search space is small.

7 Related work

7.1 Subgraph isomorphism

The *subgraph isomorphism problem* is known to be NP-complete [27]. There is an extensive literature about algorithms for special cases of the subgraph isomorphism problem, e.g., for graphs of bounded genus [41], grids [42], or planar graphs [23]. Algorithms based on searching and pruning [18,19,59] as well as constraint programming [38,70] have been used to solve the general case of the problem. Other filtering approaches that have been proposed for the subgraph isomorphism problem include filtering algorithms based on local all-different constraints [57] and reasoning on vertex dominance properties based on scoring and neighborhood functions [3]. Nevertheless, none of these approaches handle an objective function value on the isomorphic subgraph in order to find an optimal isomorphism function (i.e., *deployment*). In our node deployment problem, a mapping from application nodes to instances needs to be determined as in the subgraph isomorphism problem, but in addition the mapping must minimize the deployment cost. In fact, we are guaranteed that such a mapping exists, given that the physical graph is complete (i.e., all the instances are connected). As a result, the subgraph isomorphism problem can be trivially solved and the key issue lies in finding an isomorphic subgraph that is optimal. Overall, in this paper, we build upon the previous work on subgraph isomorphism by finding a

sequence of isomorphic subgraphs with decreasing deployment cost.

7.2 Overlay placement

Another way to look at the node deployment problem is to find a good overlay within the allocated instances. The networking community has invested significant effort in intelligently placing intermediate overlay nodes to optimize Internet routing reliability and TCP performance [30,55]. This community has also investigated node deployment in other contexts, such as proxy placement [40] and web server/cache placement [36,48]. In the database community, there have been studies in extensible definition of dissemination overlays [46], as well as operator placement for stream-processing systems [47]. In contrast to all of these previous approaches, ClouDiA focuses on optimizing the direct end-to-end network performance without changing routing infrastructure in a datacenter setting.

7.3 Virtual network embedding

Both the virtual network embedding problem [15,17,25,67] and the testbed mapping problem [54] map nodes and links in a virtual network to a substrate network taking into account various resource constraints, including CPU, network bandwidth, and permissible delays. Traditional techniques used in solving these problems cannot be applied to our public cloud scenario simply because treating the entire datacenter as a substrate network would exceed the instance sizes they can handle. Recent work provides more scalable solutions for resource allocation at the datacenter scale by greedily exploiting server locality [9,29]. However, this work does not take network latency into consideration. ClouDiA focuses on latency-sensitive applications and examines a different angle: We argue network heterogeneity is unavoidable in public clouds and therefore optimize the deployment as a cloud tenant rather than the infrastructure provider. Such role changing enables us to frame the problem as an application tuning problem and better capture optimization goals relevant to latency-sensitive applications. Also, we only need to consider instances allocated by a given tenant, which is a substantially smaller set than the entire data center. Of course, nothing precludes the methodology provided by ClouDiA being incorporated by the cloud provider upon allocation for a given tenant, as long as the provider can obtain latency-sensitivity information from the application.

7.4 Auto-tuning in the cloud

In the database community, there is a long tradition of auto-tuning approaches, with AutoAdmin [14] and COMFORT [63] as some of its seminal projects. Recently, more

attention has focused on auto-tuning in the cloud setting. Babu investigates how to tune the parameters of MapReduce programs automatically [7], while Jahani et al. [33] automatically analyze and optimize MapReduce programs with data-aware techniques. Lee et al. [39] optimizes resource allocation for data-intensive applications using a prediction engine. Conductor [65] assists public cloud tenants in finding the right set of resources to save cost. Both of the above two approaches are similar in spirit to ClouDiA. However, they focus on map-reduce style computation with high bandwidth consumption. Our work differs in that we focus on latency-sensitive applications in the cloud and develop appropriate auto-tuning techniques for this different setting.

7.5 Cloud orchestration

AWS CloudFormation [5] allows tenants to provision and manage various cloud resources together using templates. However, interconnection performance requirements cannot be specified. AWS also supports cluster placement groups and guarantees low network latency between instances within the same placement group. Only costly instance types are supported and the number of instances that can be allocated to the same placement group is restricted. HP Intelligent Management Center Virtual Application Network Manager [31] orchestrates virtual machine network connectivity to ease application deployment. Although it allows tenants to specify an “information rate” for each instance, there is no guarantee on pairwise network performance characteristics, especially network latency. Wrasse [50] provides a generic tool for cloud service providers to solve allocation problems. It does not take network latency into account.

7.6 Topology-aware distributed systems

Many recent large-scale distributed systems built for data centers are aware of network topology. Cassandra [37] and Hadoop DFS [13] both provide policies to prevent rack-correlated failure by spreading replicas across racks. DyradLINQ [68] runs rack-level partial aggregation to reduce cross-rack network traffic. Purlieus [45] explores data locality for MapReduce tasks also to save cross-rack bandwidth. Quincy [32] studies the problem of scheduling with not only locality but also fairness under a fine-grained resource sharing model. The optimizations in these previous approaches are both rack and application specific. By contrast, ClouDiA takes into account arbitrary levels of difference in mean latency between instances. In addition, ClouDiA is both more generic and more transparent to applications.

7.7 Network performance in public clouds

Public clouds have been demonstrated to suffer from latency jitter by several experimental studies [56,61]. Our previous work has proposed a general framework to make scientific applications jitter-tolerant in a cloud environment [72], allowing applications to tolerate latency spikes. However, this work does little to deal with stable differences in mean latency. Zaharia et al. [69] observed network bandwidth heterogeneity due to instance collocation in public clouds and has designed a speculative scheduling algorithm to improve response time of MapReduce tasks. Farley et al. [26] also exploit such network bandwidth as well as other types of heterogeneity in public clouds to improve performance. To the best of our knowledge, ClouDiA is the first work that experimentally observes network *latency* heterogeneity in public clouds and optimizes application performance by solving the node deployment problem.

8 Conclusions

We have shown how ClouDiA makes intelligent deployment decisions for latency-sensitive applications under heterogeneous latencies, which naturally occur in public clouds. We formulated the deployment of applications into public clouds as optimization problems and proposed techniques to speed up the search for high-quality deployment plans. We also presented how to efficiently obtain latency measurements without interference. Finally, we evaluated ClouDiA in Amazon EC2 with realistic workloads. ClouDiA is able to reduce the time-to-solution or response time of latency-sensitive applications by 15–55 %, without any changes to application code.

As future work, we plan to extend our formulation to support weighted communication graphs. Another direction of practical importance is quantifying over-allocation cost and analyzing its impact on total cost-to-solution for scientific applications. Finally, we will investigate the deployment problem under other criteria, such as bandwidth, for additional classes of cloud applications.

Acknowledgments The third author took inspiration for the title from the name of an early mentor, Claudia Bauzer Medeiros, to whom he is thankful for introducing him to database research during his undergraduate studies. We thank the anonymous PVLDB 2013 and VLDB Journal reviewers for their insightful comments which helped us to improve this paper. This research has been supported by the NSF under Grants CNS-1012593, IIS-0911036, by an AWS Grant from Amazon, by the iAd Project funded by the Research Council of Norway and by a Google Research Award. Any opinions, findings, conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and

reproduction in any medium, provided the original author(s) and the source are credited.

Appendix 1: Problem complexity

Theorem 1 *The Longest Link Node Deployment Problem (LLNDP) is NP-hard.*

Proof We reduce the *subgraph isomorphism problem* (SIP), which is known to be NP-hard [27], to the LLNDP. Consider an instance of SIP, where $G_1 = (V_1, E_1)$, $G_2 = (V_2, E_2)$, and we look for a mapping $\sigma : V_1 \rightarrow V_2$ such that whenever $(i, j) \in E_1$, then $(\sigma(i), \sigma(j)) \in E_2$. We build an instance of the LLNDP as follows. We set $N = V_1$, $S = V_2$, $G = (V_1, E_1)$, and the costs $\mathcal{C}_{\mathcal{L}}(i, j)$ to 1 whenever the edge (i, j) belongs to E_2 , and to 2 otherwise. By solving LLNDP, we get a deployment plan \mathcal{D} . G_2 contains a subgraph isomorphic to G_1 whenever $\mathcal{C}_{\mathcal{D}}^{\text{LL}} = 1$ and $\sigma = \mathcal{D}$. \square

In order to show hardness of approximation, we will assume in the next two theorems that all communication costs are distinct. This assumption is fairly realistic, even more so as these costs are typically real numbers that are experimentally measured.

Theorem 2 *There is no α -absolute approximation algorithm to the Longest Link Node Deployment Problem in the case where all communication costs are distinct, unless P=NP.*

Proof Consider an instance I of the LLNDP, consisting of $G = (N, E)$, $\mathcal{C}_{\mathcal{L}}(i, j)$ where $i, j \in S$, and $\mathcal{C}_{\mathcal{L}}(i, j) = \mathcal{C}_{\mathcal{L}}(i', j')$ if and only if $i = i'$ and $j = j'$. We order all the communication links $(i_1, j_1), (i_2, j_2) \dots, (i_{|N|^2}, j_{|N|^2})$ in increasing order of their communication costs. Let (i_w, j_w) be the longest link used in the optimal solution for I , with optimal value $\mathcal{C}_{\mathcal{L}}(i_w, j_w)$. Notice that any instance of LLNDP that shares the same G and S as well as the same exact ordering of the communication links will also have an optimal value equal to $\mathcal{C}_{\mathcal{L}}(i_w, j_w)$. Now, assume that there is an α -absolute approximation algorithm \mathcal{A} to LLNDP. We build an instance I' by changing the communication costs in I to $\mathcal{C}_{\mathcal{L}}(i_k, j_k) = (\alpha + 1)k$. Note that I and I' share the same ordering of the communication links, and any two links in I' have communication costs that are at least $\alpha + 1$ apart. Since \mathcal{A} returns a longest link $(i_{w'}, j_{w'})$ for I' such that $\mathcal{C}_{\mathcal{L}}(i_{w'}, j_{w'}) \leq \mathcal{C}_{\mathcal{L}}(i_w, j_w) + \alpha$, \mathcal{A} actually solves I' optimally. The fact that LLNDP is NP-hard completes the proof. \square

Theorem 3 *There is no ϵ -relative approximation algorithm to the Longest Link Node Deployment Problem in the case where all costs are distinct, unless P=NP.*

Proof As in the previous proof, we build an instance I' that differs from I only by the costs of the links. We set these costs to be $\mathcal{C}_{\mathcal{L}}(i_k, j_k) = (\epsilon + 1)^k$ for every link (i, j) where $i, j \in S$. In that case, for any two links (i_p, j_p) and (i_q, j_q) where $p < q$, we have $\mathcal{C}_{\mathcal{L}}(i_p, j_p) < \epsilon \cdot \mathcal{C}_{\mathcal{L}}(i_q, j_q)$. The fact that a ϵ -relative approximation algorithm would return a longest link $(i_{w'}, j_{w'})$ of I' such that $\mathcal{C}_{\mathcal{L}}(i_{w'}, j_{w'}) \leq \epsilon \cdot \mathcal{C}_{\mathcal{L}}(i_w, j_w)$ completes the proof. \square

Theorem 4 *The Longest Path Node Deployment Problem (LPNDP) is NP-hard.*

Proof The proof is otherwise identical to the proof of Theorem 1 except when (i, j) does not belong to E_2 , we set $\mathcal{C}_{\mathcal{L}}(i, j)$ to $|E_1| + 1$ and the final check is now $\mathcal{C}_{\mathcal{D}}^{\text{LP}} \leq |E_1|$. \square

Appendix 2: Distance approximations

All techniques in Sect. 5 may require non-negligible measurement time to obtain pairwise mean latencies. We also experimented with the following two approximations to network distance, which are both simple and intuitively related to mean latency.

IP Distance

Our first approximation makes use of internal IPv4 addresses in the cloud. The hypothesis is that if two instances share a common /24 address prefix, then these instances are more likely to be located in the same or in a nearby rack than if the two instances only share a common /8 prefix. We can therefore define *IP distance* as a measure of the dissimilarity between two IP addresses: Two instances sharing the same /x address prefix but not /x+1 address prefix have IP distance $32 - x$. We can future adjust the sensitive of this measurement by considering $g(1 \leq g < 32)$ consecutive bits of the IP addresses together.

Hop Count

A slightly more sophisticated approximation is the hop count between two instances. The hop count is the number of intermediate routers through which data flows from source to destination. Hop count can be obtained by sending packets from source to destination and monitoring the Time To Live (TTL) field of the IP header.

Experimental evaluation We evaluated the two approximations above with the same experimental setup described in Sect. 6.2. We compare both IP distance and hop count against the mean round-trip latency measurement results obtained using the *staged* approach described in Sect. 5.

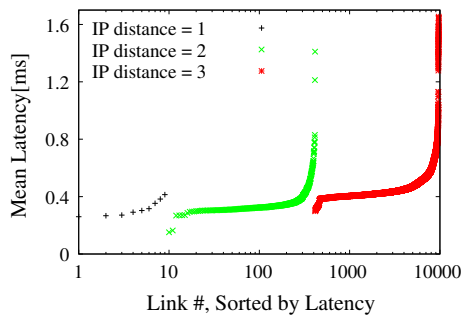


Fig. 16 Latency order by IP distance

In Fig. 16, we show the effect of using IP distances as an approximation. In this experiment, we consider 8 consecutive bits of the IP address together: two instances sharing a /24 address prefix have IP distance 1; two instances with the same /16 prefix but not /24 prefix have IP distance 2, and so on. We also experimented with other sensitivity configurations and the results are similar. The *x* axis is in log scale and divided into three groups based on the value of IP distance. Since we used the internal IP addresses of EC2, all of which share the same /8 prefix, we do not observe IP distance of 4. Within each group, links are ordered by round-trip latency measurements. If IP distance were a desirable approximation, we would expect that pairs with higher IP distance would also have higher latencies. Figure 16 shows that such monotonicity does not always hold. For example, within the group of IP distance = 2, there exist links having lower latency than some links of IP distance = 1, as well as links having higher latency than some links of IP distance = 3. Interestingly, the lowest latencies are observed in pairs with IP distance = 2.

The effect of using hop count as an approximation is shown in Fig. 17. Similarly, the *x* axis is in log scale and divided into three groups based on hop count. We do not observe any pair of instances that are two hops apart. Within each group, links are ordered by round-trip latencies. As in Fig. 16, there exists a significant number of link pairs ordered inconsistently by hop count and measured latency.

The above results demonstrate that IP distance and hop count, though easy to obtain, do not effectively predict network latency.

Appendix 3: Public cloud service providers

To demonstrate the applicability of ClouDiA in public clouds other than the one offered by Amazon Web Services, we report latency heterogeneity and mean latency stability measurements in Google Compute Engine and Rackspace Cloud Server.

Figure 18 shows the CDF of the mean pairwise end-to-end latencies among 50 Google Compute Engine n1-standard-1

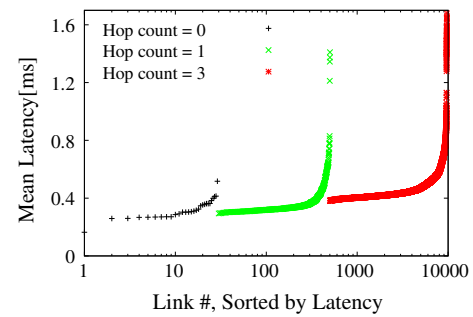


Fig. 17 Latency order by Hop count

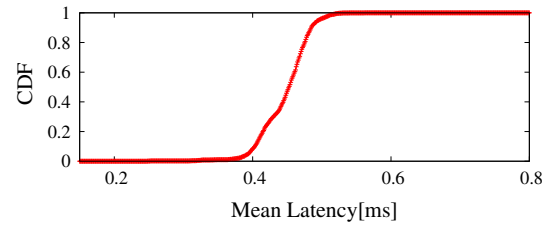


Fig. 18 Latency heterogeneity in Google Compute Engine

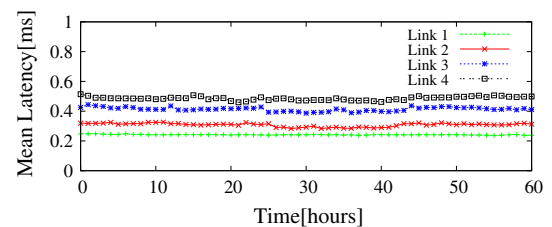


Fig. 19 Mean latency stability in Google Compute Engine

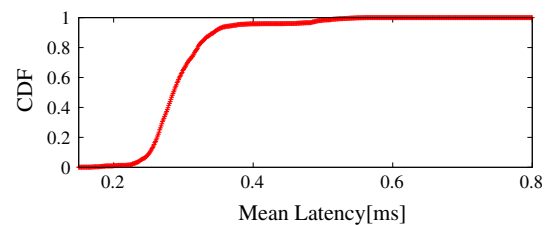


Fig. 20 Latency heterogeneity in Rackspace Cloud Server

instances in the us-central1-a region, obtained by measuring TCP round-trip times of 1 KB messages. Around 5 % of the instance pairs exhibit mean latency below 0.32 ms, whereas the top 5 % are above 0.5 ms. Figure 19 shows the mean latencies of four representative links over 60h, with each latency measurement averaged over an hour. We observe a similar behavior of mean latency stability over time as in Amazon EC2. By contrast, latency heterogeneity is somewhat smaller; however, it is still present.

Similarly, Fig. 20 shows the CDF of the mean pairwise end-to-end latencies among 50 Rackspace Cloud Server performance 1–1 instances in the Northern Virginia (IAD) region, obtained by measuring TCP round-trip times of 1 KB

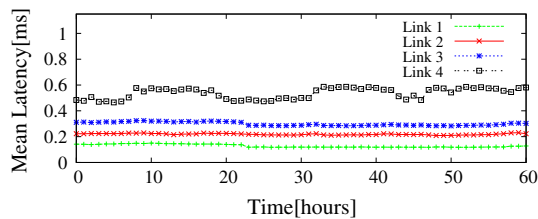


Fig. 21 Mean latency stability in Rackspace Cloud Server

messages. Around 5% of the instance pairs have latency below 0.24ms, whereas the top 5% are above 0.38ms. Figure 21 shows the mean latencies of four representative links over 60h, with latency measurement averaged over an hour. The effects observed are largely in line with the ones seen in the Google Compute Cloud.

The above results confirm the existence of latency heterogeneity and mean latency stability in the public clouds of both Google Compute Engine and Rackspace Cloud Server. The results suggest that by adopting ClouDiA, cloud tenants can achieve significant reduction in time-to-solution or service response time in these public clouds as well, and not only on Amazon E C2.

References

- Alpert, R.D., Philbin, J.F.: cBSP: zero-cost synchronization in a modified BSP model. Technical report, NEC Research Institute (1997)
- Amazon web services, elastic compute cloud (ec2). <http://aws.amazon.com/ec2/>
- Audemard, G., Lecoutre, C., Samy-Modeliar, M., Goncalves, G., Porumbel, D.: Scoring-based neighborhood dominance for the subgraph isomorphism problem. In: O'Sullivan, B. (ed.) Principles and Practice of Constraint Programming, 20th International Conference, CP 2014, 8-12 Sept 2014, Lyon, France, vol. 8656, pp. 125–141. Springer, Switzerland (2014)
- Amazon web services, case studies. <http://aws.amazon.com/solutions/case-studies/>
- Amazon web services, cloudformation. <http://aws.amazon.com/cloudformation/>
- Amazon web services, search engines & web crawlers. <http://aws.amazon.com/search-engines/>
- Babu, S.: Towards automatic optimization of MapReduce programs. In: SOCC (2010)
- Badue, C.S., Baeza-Yates, R.A., Ribeiro-Neto, B.A., Ziviani, N.: Distributed query processing using partitioned inverted files. In: SPIRE, pp.10–20 (2001)
- Ballani, H., Costa, P., Karagiannis, T., Rowstron, A.: Towards predictable datacenter networks. In: SIGCOMM (2011)
- Battre, D., Frejnik, N., Goel, S., Kao, O., Warneke, D.: Evaluation of network topology inference in opaque compute clouds through end-to-end measurements. In: IEEE CLOUD (2011)
- Benson, T., Akella, A., Maltz, D.A.: Network traffic characteristics of data centers in the wild. In: Internet Measurement Conference (2010)
- Bonorden, O., Juurlink, B., von Otte, I., Rieping, I.: The paderborn university BSP (PUB) library. *Parallel Comput.* **29**(2), 187–207 (2003)
- Borthakur, D.: The hadoop distributed file system: architecture and design. <http://hadoop.apache.org/core/docs/current/hdfsdesign.pdf>
- Chaudhuri, S., Narasayya, V.: An efficient, cost-driven index selection tool for microsoft SQL server. In: VLDB (1997)
- Cheng, X., Su, S., Zhang, Z., Wang, H., Yang, F., Luo, Y., Wang, J.: Virtual network embedding through topology-aware node ranking. *SIGCOMM CCR* **41**(2), 38–47 (2011)
- Chowdhury, M., Zaharia, M., Ma, J., Jordan, M.I., Stoica, I.: Managing data transfers in computer clusters with orchestra. In: SIGCOMM (2011)
- Chowdhury, N.M.M.K., Rahman, M.R., Boutaba, R.: Virtual network embedding with coordinated node and link mapping. In: INFOCOM (2009)
- Cordella, L.P., Foggia, P., Sansone, C., Vento, M.: Performance evaluation of the VF graph matching algorithm. In: International Conference on Image Analysis and Processing (1999)
- Cordella, L.P., Foggia, P., Sansone, C., Vento, M.: A (sub)graph isomorphism algorithm for matching large graphs. *IEEE Trans. PAMI* **26**, 1367–1372 (2004)
- Couzin, I., Krause, J., Franks, N., Levin, S.: Effective leadership and decision-making in animal groups on the move. *Nature* **433**(7025), 513–516 (2005)
- DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Voshall, P., Vogels, W.: Dynamo: amazon's highly available key-value store. In: SOSP (2007)
- Demmel, J., Hoemmen, M., Mohiyuddin, M., Yelick, K.A.: Avoiding communication in sparse matrix computations. In: IPDPS (2008)
- Eppstein, D.: Subgraph isomorphism in planar graphs and related problems. In: SODA (1995)
- Evangelinos, C., Hill, C.N.: Cloud computing for parallel scientific HPC applications. In: Cloud Computing and Its Applications (2008)
- Fajjari, I., Aitsaadi, N., Pujolle, G., Zimmermann, H.: VNE-AC: virtual network embedding algorithm based on ant colony metaheuristic. In: IEEE International Conference on Communications (2011)
- Farley, B., Varadarajan, V., Bowers, K., Juels, A., Ristenpart, T., Swift, M.: More for your money: exploiting performance heterogeneity in public clouds. In: SOCC (2012)
- Garey, M.R., Johnson, D.S.: *Computers and Intractability*. Freeman, San Francisco (1979)
- Geambasu, R., Gribble, S.D., Levy, H.M.: Cloudviews: communal data sharing in public clouds. In: HotCloud (2009)
- Guo, C., Lu, G., Wang, H.J., Yang, S., Kong, C., Sun, P., Wu, W., Zhang, Y.: Secondnet: a data center network virtualization architecture with bandwidth guarantees. In: CoNEXT (2010)
- Han, J., Watson, D., Jahanian, F.: Topology aware overlay networks. In: INFOCOM, IEEE (2005)
- HP intelligent management center virtual application network manager. http://h17007.www1.hp.com/us/en/products/network-management/IMC_VANM_Software/index.aspx
- Isard, M., Prabhakaran, V., Currey, J., Wieder, U., Talwar, K., Goldberg, A.: Quincy: fair scheduling for distributed computing clusters. In: SOSP (2009)
- Jahani, E., Cafarella, M.J., Ré, C.: Automatic optimization for mapreduce programs. In: PVLDB (2011)
- Juve, G., Deelman, E., Vahi, K., Mehta, G., Berriman, B., Berman, B.P., Maechling, P.: Scientific workflow applications on amazon EC2. In: IEEE International Conference on e-Science (2009)
- Kim, J.-S., Ha, S., Jhon, C.S.: Efficient barrier synchronization mechanism for the BSP model on message-passing architectures. In: IPPS/SPDP (1998)

36. Krishnan, P., Raz, D., Shavitt, Y.: The cache location problem. *IEEE/ACM Trans. Netw.* **8**(5), 568–582 (2000)
37. Lakshman, A., Malik, P.: Cassandra: a decentralized structured storage system. *SIGOPS OSR* **44**(2), 35–40 (2010)
38. Larrosa, J., Valiente, G.: Constraint satisfaction algorithms for graph pattern matching. *Math. Struct. Comput. Sci.* **12**(4), 403–422 (2002)
39. Lee, G., Tolia, N., Ranganathan, P., Katz, R.H.: Topology-aware resource allocation for data-intensive workloads. In: *SIGCOMM CCR* (2011)
40. Li, B., Golin, M.J., Italiano, G.F., Deng, X., Sohraby, K.: On the optimal placement of web proxies in the internet. In: *INFOCOM* (1999)
41. Miller, G.L.: Isomorphism testing for graphs of bounded genus. In: *STOC* (1980)
42. Miller, Z., Orlin, J.B.: Np-completeness for minimizing maximum edge length in grid embeddings. *J. Algorithms* **6**, 10–16 (1985)
43. Mysore, R.N., Pamboris, A., Farrington, N., Huang, N., Miri, P., Radhakrishnan, S., Subramanya, V., Vahdat, A.: Portland: a scalable fault-tolerant layer 2 data center network fabric. In: *SIGCOMM* (2009)
44. O'Hanlon, C.: A conversation with Werner Vogels. *ACM Queue* **4**(4), 14–22 (2006)
45. Palanisamy, B., Singh, A., Liu, L., Jain, B.: Purlieus: locality-aware resource allocation for mapreduce in a cloud. In: *SC* (2011)
46. Papaemmanouil, O., Ahmad, Y., Çetintemel, U., Jannotti, J., Yildirim, Y.: Extensible optimization in overlay dissemination trees. In: *SIGMOD* (2006)
47. Pietzuch, P.R., Ledlie, J., Shneidman, J., Roussopoulos, M., Welsh, M., Seltzer, M.I.: Network-aware operator placement for stream-processing systems. In: *ICDE* (2006)
48. Qiu, L., Padmanabhan, V.N., Voelker, G.M.: On the placement of web server replicas. In: *INFOCOM* (2001)
49. Rackspace. <http://www.rackspace.com/>
50. Rai, A., Bhagwan, R., Guha, S.: Generalized resource allocation for the cloud. In: *SOCC* (2012)
51. Ramakrishnan, L., Jackson, K.R., Canon, S., Cholia, S., Shalf, J.: Defining future platform requirements for e-Science clouds. In: *SOCC* (2010)
52. Ramakrishnan, R.: Data serving in the cloud. In: *LADIS* (2010)
53. Ramasubramanian, V., Malkhi, D., Kuhn, F., Balakrishnan, M., Gupta, A., Akella, A.: On the treeness of internet latency and bandwidth. In: *SIGMETRICS* (2009)
54. Ricci, R., Alfeld, C., Lepreau, J.: A solver for the network testbed mapping problem. *SIGCOMM CCR* **33**(2), 65–81 (2003)
55. Roy, S., Pucha, H., Zhang, Z., Hu, Y.C., Qiu, L.: Overlay node placement: analysis, algorithms and impact on applications. In: *ICDCS* (2007)
56. Schad, J., Dittrich, J., Quiané-Ruiz, J.-A.: Runtime measurements in the cloud: observing, analyzing, and reducing variance. In: *PVLDB*, vol. 3(1) (2010)
57. Solnon, C.: Alldifferent-based filtering for subgraph isomorphism. *Artif. Intell.* **174**(12–13), 850–864 (2010)
58. Song, Y.J., Aguilera, M.K., Kotla, R., Malkhi, D.: RPC chains: efficient client–server communication in geodistributed systems. In: *NSDI* (2009)
59. Ullmann, J.R.: An algorithm for subgraph isomorphism. *J. ACM* **23**, 31–42 (1976)
60. Vogels, W.: Data access patterns in the amazon.com technology platform. In: *VLDB*, p. 1 (2007)
61. Wang, G., Ng, T.S.E.: The impact of virtualization on network performance of amazon EC2 data center. In: *INFOCOM* (2010)
62. Wang, G., Salles, M.A.V., Sowell, B., Wang, X., Cao, T., Demers, A.J., Gehrke, J., White, W.M.: Behavioral simulations in mapreduce. In: *PVLDB* (2010)
63. Weikum, G., Hasse, C., Moenkeberg, A., Zabback, P.: The COMFORT automatic tuning project, invited project review. *Inf. Syst.* **19**(5), 381–432 (1994)
64. Wen, Y.: Scalability of dynamic traffic assignment. PhD thesis, Massachusetts Institute of Technology (2008)
65. Wieder, A., Bhatotia, P., Post, A., Rodrigues, R.: Orchestrating the deployment of computations in the cloud with conductor. In: *NSDI* (2012)
66. Windows azure. <http://www.windowsazure.com/>
67. Yu, M., Yi, Y., Rexford, J., Chiang, M.: Rethinking virtual network embedding: substrate support for path splitting and migration. *SIGCOMM CCR* **38**(2), 17–29 (2008)
68. Yu, Y., Isard, M., Fetterly, D., Budi, M., Erlingsson, Ú., Gunda, P.K., Currey, J.: DryadLINQ: a system for general-purpose distributed data-parallel computing using a high-level language. In: *OSDI* (2008)
69. Zaharia, M., Konwinski, A., Joseph, A.D., Katz, R.H., Stoica, I.: Improving mapreduce performance in heterogeneous environments. In: *OSDI* (2008)
70. Zampelli, S., Deville, Y., Solnon, C.: Solving subgraph isomorphism problems with constraint programming. *Constraints* **15**(3), 327–353 (2010)
71. Zou, T., Le Bras, R., Salles, M.V., Demers, A., Gehrke, J.: ClouDiA: a deployment advisor for public clouds. In: *PVLDB* vol. 6(2), pp. 109–120 (2012)
72. Zou, T., Wang, G., Salles, M.V., Bindel, D., Demers, A., Gehrke, J., White, W.: Making time-stepped applications tick in the cloud. In: *SOCC* (2011)