Enhanced Evaluation of the Interdomain Routing System for Balanced Routing Scalability and New Internet Architecture Deployments

Jianli Pan, *Student Member, IEEE*, Raj Jain, *Fellow, IEEE*, and Subharthi Paul, *Student Member, IEEE*

***Abstract*—Internet is facing many challenges that cannot be solved easily through ad hoc patches. To address these challenges, many research programs and projects have been initiated and many solutions are being proposed. However, before we have a new architecture that can motivate Internet service providers (ISPs) to deploy and evolve, we need to address two issues: 1) know the current status better by appropriately evaluating the existing Internet; and 2) find how various incentives and strategies will affect the deployment of the new architecture. For the first issue, we define a series of quantitative metrics that can potentially unify results from several measurement projects using different approaches and can be an intrinsic part of future Internet architecture (FIA) for monitoring and evaluation. Using these metrics, we systematically evaluate the current interdomain routing system and reveal many "autonomous-system-level" observations and key lessons for new Internet architectures. Particularly, the evaluation results reveal the imbalance underlying the interdomain routing system and how the deployment of FIAs can benefit from these findings. With these findings, for the second issue, appropriate deployment strategies of the future architecture changes can be formed with balanced incentives for both customers and ISPs. The results can be used to shape the short- and long-term goals for new architectures that are simple evolutions of the current Internet (so-called dirty-slate architectures) and to some extent to clean-slate architectures.**

***Index Terms*—Balanced incentives, future Internet architectures (FIA), interdomain routing, quantitative metrics, routing evaluation, routing scalability.**

I. INTRODUCTION

THE ORIGINAL Internet designers did not expect such a broad expansion of the Internet as it is today. It was designed almost 40 years ago for a small number of trusted communities of universities and institutions, but now, it is broadly used in business-related contexts. This trend introduces significant challenges such as routing scalability [1], mobility and multihoming, renumbering, traffic engineering, and policy enforcement [2]. Many new designs and services have to be supported via ad hoc patches. For example, mobility and security designs and services were not provided and were not supported in the original Internet architecture, and then patches such as Mobile IP [32] and IPSec [33] were incrementally designed and applied. Many of the patches are not consistent with each other, and some basic challenges cannot be solved by these patches. As a result, the Mobile IP and similar techniques have seen limited deployments.

Hence, many research programs and projects have been initiated globally targeting at changes to the current Internet architecture. We use the term "new Internet architecture (NIA)" to denote not only the clean-slate proposals but also those that are simple evolutions of the current Internet architecture. Details of future Internet research can be found in [3] and [4].

However, few NIA designs, particularly the clean-slate ones [5], consider the real incentives for multiple stakeholders, including customers and various Internet service providers (ISPs). They may also be hard to experiment, validate, evaluate, and deploy over the current Internet. The underlying reasons are complicated, but one of the most important is the existence of multiple stakeholders represented by different ISPs who may be unwilling to disclose their commerce-sensitive data [6] and also reluctant to allow experiments on their networks. Global Environment for Network Innovations [7] aims to create a cooperative infrastructure to enable research experimentations. However, for the real-world deployments, more research is still needed to study the incentives and the impact of multiple possible strategies.

Therefore, to motivate the ISPs to deploy a new architecture, we need to address two issues.

1) Know the current status better by appropriately evaluating the existing Internet, particularly the interdomain routing system.
2) Find how various incentives and strategies can affect the deployment of new architectures.

In other words, NIA designers need more evaluation approaches and results to understand the most up-to-date status before designing and implementing some key features. Such evaluation can also help the designers to form a "big picture" for appropriate deployment strategies.

However, currently, there is a "gap" between the research on NIAs and their evaluation efforts. First, new designs may ignore some practical constraints, and it is difficult to evaluate a still-not-existent architecture. Thus, the designers are eager for

Manuscript received December 15, 2012; revised July 26, 2013; accepted August 30, 2013. This work was supported in part by the National Science Foundation's Directorate for Computer and Information Science and Engineering (CISE) under Grant 1249681 and in part by Cisco University Research Program under a Grant.

The authors are with the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO 63130 USA (e-mail: jp10@cse.wustl.edu; jain@cse.wustl.edu; pauls@cse.wustl.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSYST.2013.2281129

TABLE I
SUMMARY AND COMPARISON OF NEW CONTRIBUTIONS OF THIS PAPER

	This paper	Previous work [2]
Methodology	With systematic quantitative metrics; detailed discussion and comparison over the data sources	No quantitative metrics; simple data usage and no detailed discussion
New metrics	PC, AR, CAR, APAR, RPAR, etc. (details in Section V)	None
Evaluated Items	All the metrics, and possible improvements	Only routing table size and multihoming
Deployment Incentives and Strategies	Consider various incentives, motivators, and deployment strategies (as in Table III); with systematic evaluation of each strategy on key metrics	Coarsely divided into scalability, mobility, and multihoming – driven modes; no corresponding evaluation of each mode and on each metric
Interpretation of the evaluation results	Detailed interpretation and observation of the results of each metric	Lack physical meaning discussion on the evaluation
Metric boundary value evaluation and their practical usage	Yes, consider the boundary conditions and use it to monitor and evaluate the deployment process	No discussion on boundary value and practical usage
Short- and long-term implications analysis	Yes, with detailed discussion on both short-term and long-term goals	No short-term and long-term analysis and discussion
Modeling and evaluation limitation discussion	Yes, with discussion on the evaluation constraints and limitations	No discussion on the limitations of the method used

specialized evaluations to support their assumptions, validate their ideas, and improve their solutions. Second, most existing evaluation efforts on the current Internet are not for new architectures, and most evaluation experts do not know what the architecture designers really need. Due to this gap, some key information is not available to the designers.

In this paper, we try to fill the gap through a new evaluation method. The major innovations can be described in two aspects.

1) We define several new systematic quantitative metrics and present corresponding evaluations leading to a series of new insights over the interdomain routing system through autonomous-system (AS)-level evaluation combining multiple public Border Gateway Protocol (BGP) [8] data sources. Our analysis unifies results from several measurement projects using different approaches and can potentially be an intrinsic part of future Internet architectures (FIAs) for monitoring and evaluation.
2) The new findings can help stakeholders (both customers and ISPs) to make the routing system more scalable and to formulate appropriate deployment strategies for future architectures with balanced incentives.

The rest of this paper is organized as follows. Section II describes related work. Section III presents the basic methodology. The incentives and deployment strategies are discussed in Section IV. Detailed evaluation and analysis are in Section V. Discussions on applying to new architectures are in Section VI. We conclude this paper in Section VII.

II. RELATED WORK

In this section, we discuss some related work.

A. Relationship to Our Previous Work

This paper is a significant step forward over our previous preliminary evaluation work started in [2]. We present new findings and add a series of new major contributions in the context of a new architecture design and try to form a complete evaluation framework. It also offers a new angle of view into the evaluation of the routing system for user- and ISP-level incentives and strategies. Table I briefly describes the new contributions and compares them to the previous work.

B. Debate Over Address Deaggregation of the Internet

There are different opinions on how severe the Internet address deaggregation problem is and how it contributes to the routing scalability challenge. Cittadini *et al.* [9] hold an optimistic opinion and argue that routing table size growth has not turned worse in recent years, which is opposite of the previous alarming findings [10]–[12] and the conclusion reached in [1] by the Routing Research Group (RRG) [13] of the Internet Research Task Force (IRTF). In this paper, we go beyond the debate and aim to view this issue in a longer term perspective, i.e., in the context of new architectures rather than restricted to the current status. When the IPv4 address pool finally depletes and new architectures are deployed, our goal is to find a way to make the development sustainable, the transition fluent, and the evolution transparent.

C. AS-Level Measurements

We are not the first to do AS-level analysis. In the last decade, many papers have used the public BGP data sources to analyze the Internet for different purposes such as Internet topology [14], [15], routing table size growth, and prefix aggregation [11], [16]–[18]. Roughan *et al.* [18] summarize the AS-level measurement work of the last ten years and clarify many controversial observations; the authors also argue against many “exercises-like” AS-level topology measurements. The authors advocate measuring the AS Internet as an economic construct driven by economic incentives and constrained by sociotechnological factors instead of just uninspiring abstract graphs. Our paper matches this call well.

Although some overlap with some earlier AS-level analysis is unavoidable, our approach is novel because of the following reasons: 1) we do it in a different way and for different purposes, i.e., we do it in the context of deploying new architectures and estimate how various deployment strategies can achieve different results; and 2) we set up a system of new metrics and study their boundary values for monitoring deployment progress and maintaining cost effectiveness. We

TABLE II
DATA SOURCES USED IN OUR EVALUATION AND ANALYSIS

	Route Views	CAIDA	CIDR Report
Features	Long-time logs of routing data, peering with multiple big ISPs' ASs	Provides pre-processed topology data	Multiple granularities statistical data, latest changes, summary, etc.
Comparison	Full record of a long period of time	Some results relatively old	Most up-to-date data and trends
Modeling Usage	General size estimation, trends estimation, raw data for further processing	Estimation of prefix aggregation trends (metrics APAR and RPAR), and deployment strategies and effects	General AS prefixes, BGP updates, Aggregation Ratio (AR), CAR estimation

aim to shape short-term goals such as routing scalability and long-term goals such as new architecture evolution.

D. Work on Scalability and FIA

Despite the debate over the severity of the address deaggregation [9] and the argument that hardware technology advances can make routing table size expansion issue unimportant [31], routing scalability is still one of the direct incentives (although not the only one) motivating the Internet community to look for architecture alternatives [1]. The RRG [13] of the IRTF and the Host Identity Protocol [19] and Locator/Identifier (ID) Separation Protocol (LISP) [20] groups of the Internet Engineering Task Force have been working on several solutions. A low-cost transition solution to improve the prefix aggregation is presented in [21].

Many research efforts have been recently initiated, such as Future INternet Design (FIND) [23] and latest synergistic FIA program [24] from the USA. More details of these and other global architecture design and experimentation efforts can be found in our papers [3], [4].

E. Relationship to This Work

First, new architecture ideas need more knowledge about the latest status of the Internet and about evaluation methods for deployment. Second, nontechnical incentives and deployment strategies' analysis are needed for solving the routing scalability issue and for long-term architecture evolution. Hence, these two aspects illustrate our unique goals and contributions: 1) to provide systematic evaluations on the real-world status for deployment reference; and 2) to provide incentives and deployment strategies' analysis for NIAs.

III. METHODOLOGY

Internet is too complex, and there is not a single place that can claim having a complete copy of raw data characterizing the "whole picture" of the Internet. Hence, to avoid or limit the chances of biased information, we combine several public interdomain routing data sources. The three data sources we used include Cooperative Association for Internet Data Analysis [25], Route Views [26], and Classless InterDomain Routing (CIDR) Report [11]. They have different features and we use them for different purposes. We also carry out cross comparisons among these sources to validate the results and observations. For incentives and deployment strategies' analysis, we combine the data sources to get a clearer picture of the many factors that can impact any new architectural deployment. We evaluate how the different incentives and deployment strategies lead to different results in achieving short- and long-term goals. The features, comparison, and modeling usage of the three data sources are listed in Table II.

Our evaluation is based on AS-level analysis. Although there are limitations of such methods for other measurement purposes [18], we find it useful in our evaluation for new architecture research. Although AS is always deemed as only a basic routing unit, it is also a basic organizational unit, which is overloaded with both interdomain routing policy and high-level business relationships [2] (indicating cash flow of customer-to-provider and provider-to-provider relationships). These later parts and the real incentives for the stakeholders are always overlooked. Much information is hidden among the interaction of the ASs, and we want to reveal it through quantitative evaluation. We believe that such methods can reveal significant organizational information for the future Internet. These concepts about ASs were also reflected in our Mobility and Multihoming supporting Identifier Locator Split Architecture (MILSA) architecture [2] and policy-oriented frameworks [27].

In this paper, we define unified quantitative metrics to evaluate the current status. Some evaluations have used different approaches on different data sources and may have drawn biased conclusions. Using a unified metrics can help mitigate such errors and form a complete view of the current Internet by integrating different measurements. In addition, for new architectures, a quantitative metrics system integrated into the architecture is necessary to monitor and control the Internet. Our effort can be a good starting point for this. Moreover, we try to avoid the mistakes made by several previous BGP data analysis papers that only look at the data without identifying the structural components that cause the observations. Hence, in the evaluation for each metric, in addition to analyzing the data, we have observations and key lesson discussions in the context of NIAs, which is one of the key goals of this paper.

IV. FUTURE INTERNET: DEPLOYMENT INCENTIVES AND STRATEGIES

In view of the debate regarding the future architecture to be a clean-slate or an evolutionary one [28], in this paper, we focus on evolutionary ones since we believe that today's Internet is too big to be started over again.

A. Deployment Incentives

We do a detailed evaluation in the next section to address the problems faced in the context of both routing scalability and future Internet evolution. However, every change needs

TABLE III
DEPLOYMENT STRATEGIES: THEIR INCENTIVES, MAJOR MOTIVATORS, AND DEPLOYMENT ORDER

	Top-down	Bottom-up	Middle-way	Adaptive
Incentives	Mobility	Multihoming, Traffic Eng., Renumbering	Scalability	Composite
Motivators	Mobile ISPs, especially big Mobile ISPs	Stub Customer ASs	Medium, Small ISPs who need investment protection most	ASs that have highest scores based on an evaluation standard
Deployment Order	Big ISPs → Medium, Small ISPs → Stub ASs	Stub ASs → Medium, Small ISPs → Big ISPs	Medium, Small ISPs → Big ISPs or Stub ASs	ASs with High Scores → ASs with low Score

balanced incentives and motivations for ISPs and customers. This also applies to the transition to any potential new architecture. For deployment, it is necessary to have a clear strategy that balances the interests of all parties and sets up a "win–win" business model in the technical deployment plan. After a thorough investigation, we summarize the incentives into three categories: 1) privacy and ***mobility*** that the new solution can provide for a single user or a group of users; 2) ***multihoming, traffic engineering,*** and ***easy renumbering***; and 3) ***scalability***. Users or ISPs may have different priorities and preferences in achieving these under certain cost constraints. End users expect to pay only for the features they need instead of all available features. ISPs are more concerned with protecting their investment and being as efficient as possible. For example, the broad usage of provider-independent addresses [1] for stub ASs (for portability, multihoming, and load balancing, etc.) put high burden on the BGP routers in terms of routing scalability. The ISPs have to keep upgrading their core devices to keep the network running. In general, they may prefer to accept solutions that are backward compatible.

In short, we need to balance such incentives and use appropriate strategies to motivate the deployment. Although there are many new architecture ideas available, LISP [22] looks closely to the real deployment. Hence, we use it as an example for some evaluations in this paper. However, our approach and findings are not restricted to LISP and are applicable to any potential NIA.

B. Deployment Strategies

Customer ASs and ISP ASs play different roles in the routing system, which means that various deployment orders can lead to different results. Specifically, with the defined metrics and the evaluation results (see Section V), we can devise multiple cost-effective deployment strategies that fit the demands and incentives of end users and ISPs best. These are the sample conceptual deployment strategies.

1) ***Bottom-up Strategy***: First deploying functional devices of the new solution for the stub customers' ASs, then the small- and medium-sized ISPs, and finally at the big ISPs. Multihoming, traffic engineering, and renumbering are the major incentives. Stub customers' ASs may prefer this strategy due to their motivation of improving these services.
2) ***Top-down Strategy***: First deploying the new solution at the edge of the big ISPs, then the medium- and small-sized ISPs, and finally at the stub customers' ASs. Mobility is the major incentive since both end user and big mobile service ISPs have demands for this. Mobile customers and big ISPs providing mobility services may prefer this strategy.

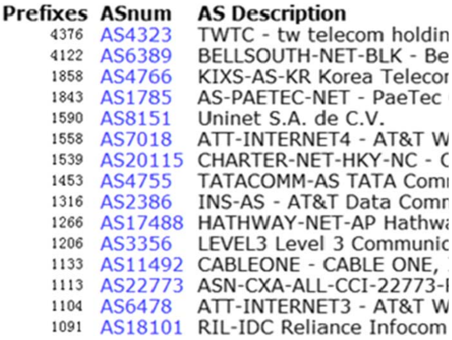

Prefixes	ASnum	AS Description
4376	AS4323	TWTC - tw telecom holdin
4122	AS6389	BELLSOUTH-NET-BLK - Be
1858	AS4766	KIXS-AS-KR Korea Telecor
1843	AS1785	AS-PAETEC-NET - PaeTec
1590	AS8151	Uninet S.A. de C.V.
1558	AS7018	ATT-INTERNET4 - AT&T W
1539	AS20115	CHARTER-NET-HKY-NC - C
1453	AS4755	TATACOMM-AS TATA Comi
1316	AS2386	INS-AS - AT&T Data Comr
1266	AS17488	HATHWAY-NET-AP Hathwa
1206	AS3356	LEVEL3 Level 3 Communic
1133	AS11492	CABLEONE - CABLE ONE,
1113	AS22773	ASN-CXA-ALL-CCI-22773-
1104	AS6478	ATT-INTERNET3 - AT&T W
1091	AS18101	RIL-IDC Reliance Infocom

Fig. 1. Top 15 ASs announcing the most prefixes.

3) ***Middle-way Strategy***: First deploying for the medium- and small-sized ISPs and then toward the two ends. The major incentive is routing scalability because medium- and small-sized ISPs have more motivation to protect their current investment and may not be able to or be willing to catch up the pace of hardware upgrades.
4) ***Adaptive Strategy***: The deployment priorities are based on a selected set of metric values and criteria. The incentive is a balanced combination of different factors. The deployment order is decided by balancing the priorities of all ASs.

We summarize the enhanced strategies and incentives, the major motivators, and the deployment priorities in Table III. We further evaluate how these strategies impact the Internet.

V. EVALUATION AND ANALYSIS

This section covers: 1) quantitative metrics definition and detailed evaluation; 2) evaluation of various deployment strategies; and 3) metrics' boundary values' evaluation.

A. Quantitative Evaluation Using New Metrics

We start from simple ones and move to the complex ones.

- **Prefix Contribution (PC):** This is the number of prefixes contributed by the AS to the total prefix entries of the "global routing table" (also known as the BGP interdomain routing table). We will use "routing table" or "table" as abbreviations for the rest of this paper.
  Metric function: This metric illustrates the overall contribution of an AS to the total routing table size.

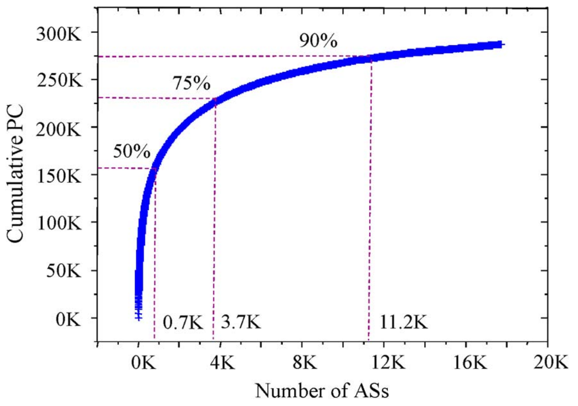

Fig. 2. Cumulative PC distribution in the whole AS space (ASs sorted in increasing order of PCs).

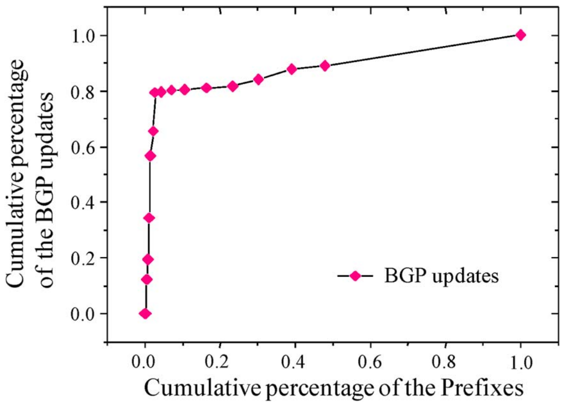

Fig. 3. One-hour sample of BGP updates in August 2009.

Evaluation: In Fig. 1, we present a list of top 15 ASs having highest prefix contributions (PCs) based on a daily-based snapshot from CIDR Report [11] in 2010. We observe that their PCs are much higher than those of the rest. For example, the AS in the first place has a PC of 4376, whereas the average PC is 9 for all ASs.

We also evaluate the cumulative PC distribution of all ASs in Fig. 2 (a data snapshot from 2006). The top 2% (700 ASs) of the total AS space contribute over 50% prefixes, whereas the top 33% ASs contribute over 90% of the total prefixes in the global routing table. The distribution reveals a significant imbalance among ASs, which is also close to the so-called "***80–20 rule***" (or Pareto principle). To verify such imbalance, we further sample a set of prefixes and the BGP update messages that they generate in 1 h in August 2009. (Note that it should not be inferred that these prefixes are from the top ASs that announce prefixes most.) As shown in Fig. 3, we observe that less than 5% of the prefixes generate around 80% of the BGP update messages. It roughly matches the above "80–20 rule" and the trend shown in Fig. 2. More on BGP update dynamics analysis and their relationships with prefixes are in [9].

Observations and Key Lessons for NIA: *From PC distribution, we can see two types of imbalance in the current interdomain routing system: 1) between ISPs' ASs and customer ASs; and 2) among ISPs' ASs. It is partially due to the structure of the current interdomain routing system and because different ASs have different roles in serving other ASs. It means that there is still plenty of room for future improvements. Deploying transitional solution with higher priorities in the ASs with higher PCs may be a good start. For NIAs, the PCs (or their counterpart in the NIA) for an AS should be proportional to the number of customers or actual usage under their territory. It also needs to be monitorable and controllable from architecture-level supports. Devising such standard metric and allowing comparison among different NIA deployments can potentially help see and evaluate the performance of NIAs.*

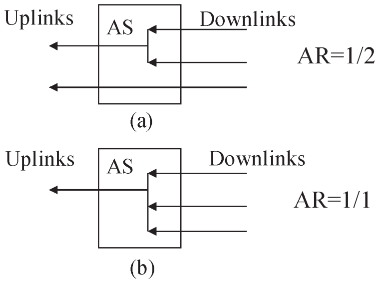

Fig. 4. AR examples.

- **Aggregation Ratio (AR):** This is the ratio of the number of prefixes announced outside after aggregation inside the AS to the total prefixes announced outside. It is described as

$$\boldsymbol{aggregation\ ratio\ (AR) = P_a/(P_a+N)} \quad \boldsymbol{(a)}.$$

Metric function: This metric approximately describes how well a specific AS performs aggregation inside itself.

P_a is the number of prefixes announced outside after aggregation inside the AS, and N represents the prefixes the AS gets from its "customer cone" that are not aggregated. (Note: For a transit AS, all its customer ASs and customer-of-customer ASs form a "cone," including the transit AS itself, and we use algorithms and data from [25] for such "cone" evaluation.) $(P_a + N)$ is the number of total prefixes announced by the AS. However, note that the prefixes are considered "aggregated inside the AS" only when there is a precise match of AS path so that the traffic transit policies are preserved. In addition, note that the aggregation does not mean perfect or maximum one. We give two examples to illustrate the AR definition and the way to compute it in Fig. 4. For example (a), $P_a = 1$ and $N = 1$; hence, AR $= 1/2$. For example (b), $P_a = 1$ and $N = 0$; hence, AR $= 1$. Note the difference between our definition and the deaggregation factor defined in [9], we focus on the aggregation behavior of an AS along the AS paths instead of simply using the ratio between the number of announced prefixes and allocated blocks.

Evaluation: We evaluate the ARs of the 32 125 ASs and sort them in decreasing order by their PCs. The results are shown in Fig. 5. Interestingly, we observe the symmetric distribution near the horizontal line of AR $= 0.44$. The top group with ARs close to 1 represents the stub customer ASs, and the dots below it with ARs ranging from 0 to 0.95 are mostly ISPs. We further divide the ISPs into two categories based on the value of AR $= 0.44$, and hence, we have three groups of ARs. Roughly, Group 1 has lowest ARs and are mostly big- and medium-sized ISPs; Group 2 consists of mostly medium- and small-sized ISPs who have AR values in the middle; and Group 3 are mostly stub customer ASs having ARs close to 1.

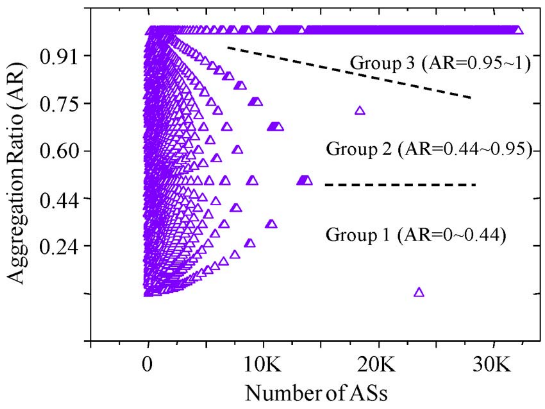

Fig. 5. AR distribution in the whole AS space (ASs are sorted in decreasing order by PCs).

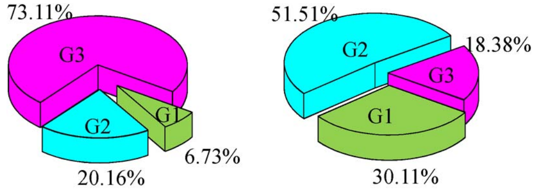

(a) (b)

Fig. 6. Number of ASs and prefixes announced in groups 1, 2, and 3.

We then study the numbers of ASs in each group and how much they contribute to the total routing table size. The results are shown in Fig. 6. Groups 1 and 2 (mostly ISPs) only cover 27% of total AS space but contribute more than 81% of total prefixes in the routing table. Moreover, we observe that big- and medium-sized ISPs generally have low ARs, and small ISPs and stub customer ASs have higher ARs due to their location near the edge. In the ideal case, if perfect aggregation is performed, every AS should have AR = 1. However, in the current Internet, the top 50 ASs with the highest PCs have an average AR of 0.44, and taking all the stub ASs' ARs into account, the average AR of the whole AS space is only 0.61.

Observations and Key Lessons for NIAs: *The aggregation for the current Internet, particularly those big ISP ASs, is far away from the ideal case due to many factors. Business relationships intermixed with routing protocol may be one of them. Although some hold positive view over the aggregation issue [9], for NIA consideration, particularly in near term when IPv6 has to come and in longer term of the world of "Internet of Things" [30] where smart devices are everywhere, the huge address space and Internet scale make better aggregation even more important. Separation of multiple-level ID spaces for dedicated functions may avoid semantic overloading and make the aggregation easier and more manageable.*

- **Cumulative Aggregation Ratio (CAR):** Of a collection of ordered ASs is the ratio of the total aggregated prefixes to the total prefixes in the table contributed by the collection.

  Metric function: This approximately describes how well the collection of ASs performs aggregation inside the collection. Detailed cumulative AR (CAR) evaluation and improvements for different strategies are presented in Section V-B3.

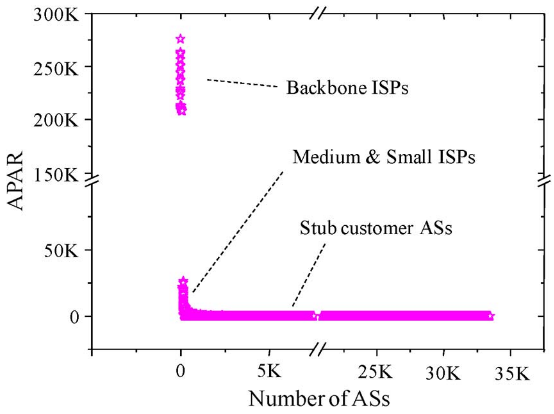

Fig. 7. APAR distribution in the whole AS space (ASs are sorted in decreasing order by APARs).

The above metrics are about a single or a group of ASs and do not reveal the interaction among ASs, particularly among ASs in the customer cone. Hence, we have more metrics.

- **Absolute Prefix-to-AS Ratio (APAR):** This is the total prefixes announced by the ASs in the customer cone of the AS. It can be expressed as

  $$\boldsymbol{APAR} = \boldsymbol{\Sigma}(\boldsymbol{PC})_{\text{cone}} \quad (\boldsymbol{b}).$$

  Metric function: The metric approximately describes the position of the AS in the AS hierarchy and how well its customer cone performs aggregation.

Evaluation: Note the difference with PC. Absolute prefix-to-AS ratio (APAR) is the sum of the PCs in the customer cone. The higher level the AS is located in the AS hierarchy, the bigger the APAR it generally has. For example, in our data sample, AS 1239 (Sprint) has APAR of 263 628, whereas AS 2552 (Washington University) has APAR of 62. We also study the APAR distribution among the whole AS space, which is shown in Fig. 7. There are very significant differences among the APARs of the three categories (i.e., backbone, middle, and stub).

Observations and Key Lessons for NIAs: *Compared with AR, the metric of APAR is a more straightforward and accurate reflection of the target ASs' position in the AS tree. Bigger APARs generally reflect poorer aggregation in the customer cone. However, they cannot reflect the exact impact of the customer cone to the target AS and need other metrics. For example, for two ASs with the same APAR, we cannot easily tell which one does aggregation better. We use relative prefix-to-AS ratio (RPAR) to help depict the inside aggregation status. For a NIA such as MILSA [2], an AS is also an organizational business management unit besides being a basic routing unit. The business-level policy is separate from routing and packet forwarding. The APAR distribution can be a lot more balanced. By dedicated control and management planes, an AS will have stronger interactive capability with its customer cone to keep the overall routing balanced and scalable.*

- **Relative Prefix-to-AS Ratio (RPAR):** This is the ratio of the total unique prefixes announced by the AS's customer

TABLE IV
APAR AND RPAR FOR THE TOP 15 ASs WITH HIGHEST APAR

	AS4323	AS6389	AS4766	AS1785	AS17488	AS7018	AS8151	AS20115	AS6478	AS2386	AS4755	AS3356	AS11492	AS22773	AS9583
APAR	4313	4178	1836	1736	1553	1497	1477	1472	1396	1289	1232	1220	1120	1095	1095
RPAR	9.44	30.77	9.42	9.37	1553	9.61	17.08	11.67	1396	147	9.42	9.45	1120	9.35	47.77

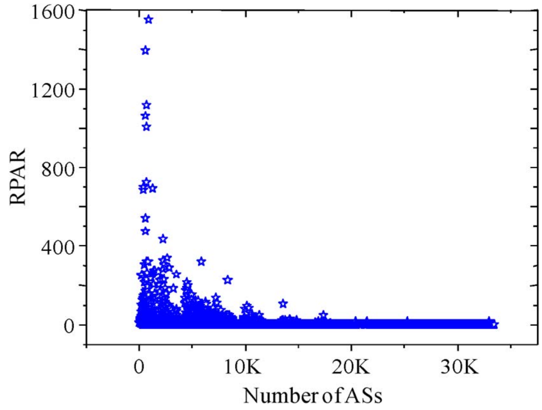

Fig. 8. RPAR distribution in the whole AS space (ASs are sorted in decreasing order by APAR).

cone to the number of total ASs in the cone

$$\boldsymbol{RPAR} = \boldsymbol{\Sigma}(\boldsymbol{PC})_{\text{cone}}/\boldsymbol{\Sigma}(\boldsymbol{AS})_{\text{cone}} \quad (\boldsymbol{c}).$$

Metric function: This approximately depicts the average aggregation status in the customer cone.

Evaluation: Intuitively, smaller RPAR means better average aggregation in that customer cone. For example, AS 1239 (Sprint) has RPAR of 9.4 and AS 2552 (Washington University) has RPAR of 62. Washington University is announcing prefixes more than the average. However, we may draw wrong conclusion if we see RPAR only. We know that Sprint is a big ISP and its customer cone included a lot of downstream ASs; hence, the RPAR of 9.428 is an average number. To explore the inner relationship between APAR and RPAR, we take the top 15 ASs in Fig. 1 and compare their APARs and RPARs. The results are shown in Table IV. It shows no absolute linear relationship between APARs and RPARs. Instead, there is a very obvious variation among the RPARs of different ASs. The variation means that even for two ASs with the same APAR, their RPARs can be significantly different due to the complexity inside their customer cone. We further study the RPAR distribution among the whole AS space, as shown in Fig. 8. Compared with APAR, it exhibits a little fluctuation for the ISPs' part, which is mostly due to the complexity inside ISP ASs.

Observations and Key Lessons for NIAs: *We may combine the APAR and RPAR to decide which ASs have better aggregation even when they have the same APAR. We may* ***take both metrics into consideration and effectively develop a deployment strategy****. Specifically, for a NIA deployment, a recursive method can be taken from bottom to top level or reversely to inspect the metric values for the ASs to decide suitable deployment strategies. In addition, the above method will help identify where the imbalance is located in the customer cone and keep the routing plane stable.*

We summarize the key metrics and their major functions in Table V. For the rest of this paper, we use the above new metrics to carry out further in-depth evaluation.

B. Deployment Benefits of Various Strategies

We focus on evaluating the effects of the strategies in reducing the routing table size, improving the PCs of each type of ASs, and improving CARs.

1) Benefit of Total Prefix Numbers' Decrease: First, we present an enhanced evaluation on the possible benefits of decreasing the routing table size. In this paper, we address three key limitations of our previous work [2]: 1) It assumed static starting point and neglected the 20% natural increase each year; 2) it considered only the scalability-driven model and only the first half phase, and no evaluation on complete scenarios of all the strategies; and 3) the predefined deployment rates (i.e., 10%, 20%, and 30% each year) were too conservative for deployment. Therefore, we address the problems and conduct a more complete evaluation.

We assume a "phase-by-phase" deployment pattern in which a specific group of ASs (such as big ISP ASs) finishes deployment following the others (such as medium and small ISP ASs or stub ASs). Hence, the deployment process is divided into two phases for all the three strategies. For the sample under study, the total routing table size starts around 310 000. Due to the deployment, the regular 20% natural increase every year will gradually decrease, and therefore, we simulated this effect in Fig. 9. Initially, the total routing table size increases but soon begins gradually decreasing due to the continuous deployment until it reaches a level close to the lower bound of the phase. In each phase, we curve the shapes to simulate the effects that the new technique deployment is relatively slow at the beginning and end periods and faster in the middle period.

Observations and Key Lessons for NIAs: Fig. 9 shows that in terms of reducing table size, *the "middle-way strategy" is more effective compared with the other two strategies in the studied period, whereas "top-down strategy" and "bottom-up strategy" are relatively close. For even longer term NIA considerations, we need to guarantee that whichever strategy is taken, finally, it will reach the same status as expected in the NIA design. The findings in this paper will help achieve cost effectiveness for the deployment process.*

The above observations validate our evaluation in Section V-A and the strategy discussion in Section IV in that the middle-way strategy is mostly driven by the large number

TABLE V

SUMMARY OF THE KEY METRICS AND THE MAJOR USAGE COMPARISON

Metrics	Metrics on individual AS's fact and behavior			Metrics reveal interactions among ASs	
	PC	AR	CAR	APAR	RPAR
Definition	Number of prefixes announced	Prefixes announced outside after aggregation inside divided by the total prefixes announced outside	Total aggregated prefixes divided by the total prefixes in the routing table for a set of ordered ASs.	Total prefixes announced by a cone of ASs	APAR divided by the number of ASs in the customer cone
Function	Depicts overall contribution of an AS to the scalability issue	Depicts how well an AS's aggregation is	Depicts how well a collection of ordered ASs' aggregation is	Depicts the position of the AS and its aggregation status	Depicts the average aggregation in the cone
Indication/ relation to other metrics	Improving aggregation of the ASs with higher PCs can improve scalability	Current aggregation is far away from ideal case of perfect aggregation with AR=1	Reveals information about a group of ordered ASs but is not able to reflect the correlation among ASs	More accurate than AR in indicating the position of AS in the Internet hierarchy	Combining APAR and RPAR can reveal more info. and help deployment

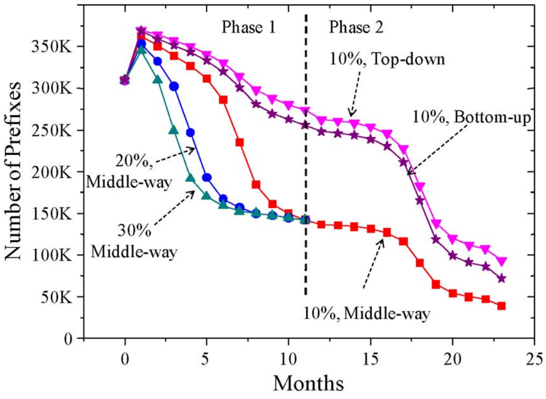

Fig. 9. Prefix decrease for the three strategies with deployment speeds of 10%, 20%, and 30% per month.

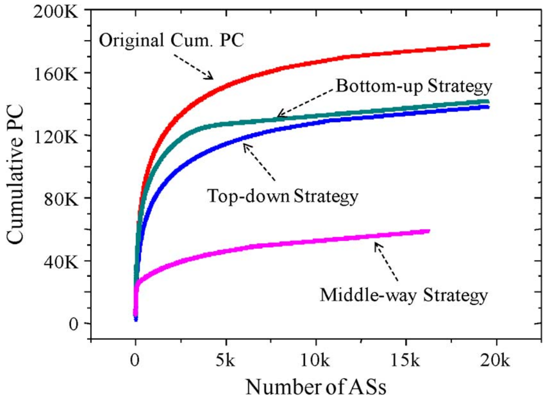

Fig. 10. Cumulative PC improvements with the three deployment strategies.

of medium and small ISPs, which are the major motivators for scalability-oriented solutions and strategies. Specifically for middle-way strategy, we consider three deployment speeds, i.e., 10%, 20%, and 30% each month; they finally meet at the endpoint of phase 1 and by then we can achieve about 44% reduction in a period of 11 months compared with about 20% reduction for the top-down and bottom-up strategies. For phase 2, the middle-way strategy finishes deployment on the big ISP and stub ASs with further 70% reduction (using the end of phase 1 as a starting point), and the top-down and bottom-up strategies finish the deployment on the small and medium ISPs with further 67% reduction in the following period of 11 months. Note that at the end of phase 2, the curves of the three strategies do not meet because phase 3 deployment is still pending. They will finally join when all phases are finished.

2) PC Improvement of Each Strategy: We further evaluate the PC improvements of the three strategies. First, we have the original cumulative PC data, as shown in Fig. 2, and we study the three strategies separately and reveal the potential difference. The results are shown in Fig. 10. We evaluate and compare the data at the end of phase 1 deployment of the three strategies. We do not consider the deployment speeds here to avoid too much information mixed altogether. Note that we sorted the data in a decreasing order to calculate cumulative PC. In Fig. 10, similar to Fig. 2, the original cumulative PC distribution shows very imbalanced status. In Fig. 10, we can see that the middle-way strategy achieves significant improvement compared with others. The shape of the curve shows that the values of PCs are smaller than the original ones. Top-down and bottom-up strategies show close effects but not as good as the middle-way strategy.

Observations and Key Lessons for NIAs: Intuitively, *the lower the curve, the better balance is achieved in the overall aggregation of the routing system, hence better scalability during the deployment progress. The difference will diminish in the NIA when the whole deployment is finished.*

3) CAR Improvements: Here, we evaluate the CAR improvements by the three strategies. We continue our preliminary evaluation in [1] by: 1) basing our evaluation on new quantitative metric of CAR; 2) adding the new middle-way strategy evaluation; and 3) adding more results and discussions on the implications.

We show the CAR improvement results in Fig. 11. Before the deployment, the CAR ranges from 0.44 for the top 50 ASs to 0.61 for the whole AS space. We sort the ASs according to their PCs in a decreasing order. For top-down strategy, we deploy first at the top ISPs, and for the first step, we assume half of their unaggregated prefixes to be aggregated. We make this assumption since we cannot expect that all the ISPs can finish the deployment overnight, but we can first aggregate the prefixes that failed to aggregate previously, hence increasing the ARs. Note that the results in this paper are based on the classification of new strategies, and we also added the curve for the "middle-way" strategy. It shows that the CARs are significantly improved under the top-down strategy. Note that there are several "turning points" on the curve, which are marked with dashed circles in Fig. 11. They happen to be at the intersections between different types of ASs, which matches the grouping results that we found in Fig. 6. For example, for the

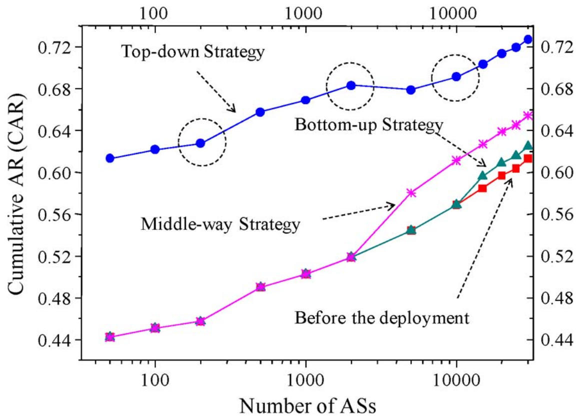

Fig. 11. CAR improvements of the three strategies [Note: 1) ASs are sorted in decreasing order by their PCs; and 2) assume a first step of 50% reduction of the unaggregated prefixes].

second point, the CAR curve turns down a little bit since most of the stub ASs cannot benefit from the top-down strategy.

Similarly, we have curves for the bottom-up and middle-way strategies. The middle-way strategy shows CARs' improvement lower than the top-down strategy but higher than the bottom-up strategy. The bottom-up strategy provides little CARs' improvement, indicating that it may not be as effective as the other two in terms of reducing the routing table size.

Intuitively, the results in Fig. 11 show that the top-down strategy improves the CARs most. However, it may be a little misleading since it does not directly match Figs. 9 and 10 in which the middle-way strategy literally shows the best effects in achieving PCs and total table size reduction. Note that in Fig. 11, the y-axis is the ***cumulative distribution function*** (cdf) of the ARs, and the x-axis is the ASs in ***logarithmic*** decreasing order sorted by their PCs. Unlike the evaluations in Figs. 9 and 10, here, we only focus on CAR. The left parts of the curves for the three strategies are mostly the tier-1 and big ISPs' ASs that have higher PCs. Therefore, deploying new solution in these ASs leads to a significant CAR improvement for these ASs since the other ASs are not counted into the cdf yet. Combining this with the above "turning points" observation, we have the following.

Observations and Key Lessons for NIAs: For aggregation improvement, *the most effective way is to follow the top-down strategy and deploy in big or tier-1 ASs first, and then combine it with other strategies by recursively improving the aggregation level by level starting from the bottom stub ASs. However, due to the overloaded semantic of current AS [2], changing the top of the AS hierarchy may need synergy from the customer cone. It is also the reason we define APAR and RPAR to reflect such interactions. For NIAs, such interactive mechanisms should be provided from architecture level such as the inter-realm communication mechanism in [2].*

C. Monitoring Deployment by Metrics' Boundary Values

We now estimate the metrics' "upper or lower bound" values in a well-aggregated Internet. Doing this can be helpful in ***monitoring the deployment progress and help put the investment into the most cost-effective parts first***. Here, the ideal case means that IP semantic overloading problem has been solved and locator aggregation is performed effectively.

- **PC:** The PC of each AS is determined by its position in the AS hierarchy. For tier-1 ASs, their PCs are at least 1 (sufficient), or more than 1 (more than minimum, but may be allowed for transition or temporary engineering purposes). Other ASs' prefixes are not needed to appear in the routing table since they will be aggregated by the tier-1 ASs' prefixes. Hence

$$\boldsymbol{PC = 1\ for\ tier\text{-}1\ ASs\ and\ 0\ for\ all\ the\ other\ ASs.} \tag{1}$$

Certainly, during the interim period before reaching the ideal case, the PC of the ASs will be larger than the values in (1) and it will get closer as the deployment proceeds.

- **AR:** In the boundary case, each AS will aggregate the prefixes under its domain and announce outside only one aggregated prefix covering all its space. This is to say that

$$\boldsymbol{AR = 1\ for\ all\ ASs.} \tag{2}$$

Note that AR = 1 does not guarantee perfect aggregation, but perfect aggregation always means AR = 1. For aggregation improvement progress, AR is a good deployment touchstone.

- **CAR:** Similar to AR, as long as the ASs perform efficient aggregation, for ASs in each level, no matter where they are located in the AS hierarchy, the CAR will be one. This is to say that

$$\boldsymbol{CAR = 1\ for\ all\ ASs.} \tag{3}$$

- **APAR:** From (1), for tier-1 ASs that have prefixes announced in the routing table, due to the perfect aggregation, their APARs will be equal to the PCs since their downstream ASs' prefixes are aggregated by their single prefix. This is to say that

$$\boldsymbol{APAR = 1\ for\ tier\text{-}1\ ASs\ and\ 0\ for\ all\ the\ other\ ASs.} \tag{4}$$

Although APAR and PC demonstrate similar boundary values, they have different meanings in evaluating the progress of deployment. Specifically, PC describes a single AS's behavior, whereas APAR shows a set of ASs' behavior in the customer cone. They cannot replace each other. Instead, combining these two is an efficient way to monitor the progress of the aggregation inside an AS customer cone.

- **RPAR:** RPAR value is the APAR divided by the number of ASs in the cone. Hence

$$\boldsymbol{RPAR = 1/\Sigma\ (AS)_{\text{cone}}\ for\ tier\text{-}1\ ASs\ and\ 0\ for\ all\ the\ other\ ASs.} \tag{5}$$

It is easy to see that the boundary RPARs for the tier-1 ASs are very close to 0, whereas the others are all 0. RPAR shares little difference between top-tier and the lower tier ASs. However, in the deployment process, the dynamics of such differences can be used as an effective way to monitor the progress.

For the total routing table size, in an ideal aggregation case, it can be the number of tier-1 ASs. However, it is also possible that each tier-1 AS may announce a little more than 1 prefix. This is controllable due to the small amount of tier-1 ASs. It will be also smaller than the total number of ASs since most lower tier ASs' prefixes are aggregated and will not appear in the interdomain routing table. Hence, the changing dynamics of the table size can be a vivid "index point" for the deployment progress of the new architecture.

VI. REASONING TO NIAs

In this section, we present the discussion on the reasoning of the evaluation results to NIAs.

A. Evaluation Results

1) Quantitative Metrics and the Corresponding Evaluations Reveal New Findings Useful for NIAs: This paper shows many facts and trends on the status of the current interdomain routing system, the contributions of different types of ASs (big, medium, and small ISPs' ASs and customer ASs), their relationships and interactions (AR, CAR, APAR, RPAR, etc.) leading to the problem space, and the underlying implications. It provides useful tips for how the problems can be effectively alleviated or solved. It also lays the foundation for further incentives and deployment strategies' evaluation.

2) Successful Design and Deployment of New Architectures Need Full Consideration on Incentives for Both Customer and ISPs and With Balanced Strategies: Any change to the Internet needs appropriate incentives to get a chance to succeed. This observation is validated by the National Science Foundation report [29] on the latest trends in the future Internet research. In the evaluation, we find that various deployment orders lead to significantly different results. Thus, we are interested in how new solutions can benefit from the findings for different interest groups. We discuss real incentives, specify their major motivators, and show how they can influence the formulation of practical deployment strategies. Further evaluation of the major benefits of the three strategies shows effectiveness and underlying implications concerning all possible improvements.

3) Boundary Value Analysis for Various Metrics Provides Good Guidelines for Monitoring and Evaluating the Deployment Progress of New Architectures: New architecture designs, even with solid technical qualities, may have to experience long deployment and evolution process. For various strategies, this process can be monitored and measured by boundary values of various metrics and can serve as an "index value" for improvements, hence achieving cost effectiveness by dynamic adjustments of these strategies.

B. Balance Issue

It is impossible to get absolute balance for the routing system since ISP ASs have different roles than stub customer ASs in the AS hierarchy. Instead, by better balanced Internet, we mean that ASs with similar roles ("type-2 imbalance," as discussed in Section V-A) should be equal in sharing the responsibilities of keeping the prefixes aggregated in a sustainable and balanced way. We can discuss the imbalance in technical and nontechnical senses. Technically, the current routing design is old, and its enforcement does not address the prefixes allocation and aggregation very well. Hence, future Internet has to address the weakness and pitfalls of the current design. With new principles, NIAs have to be open and consider all the current limitations and define clear and feasible changing steps. From nontechnical aspects, the Internet is currently intertwined with commercial interests of ISPs and customers. ***The imbalance of the routing system is an embodiment of the imbalance underlying the social commercial interests***. Hence, in the future design, the commercial and human factors need to be included into the architecture design, i.e., the future Internet is a "***network of the people***." Hence, balancing incentives for all stakeholders at every step of the changes is important.

C. Short- and Long-Term Goals of Future Internet

1) On Short-Term Goals and Effects: The evaluation provides hints on how to achieve short-term routing scalability with various speeds. For example, several quantitative metrics' values can be improved directly or indirectly. For direct improvement of an individual AS, we may first improve the PC and AR since improving the inner aggregation of each AS is relatively easier than improving a whole AS customer cone. By carrying out such improvements step by step in different ISPs' ASs, the overall routing scalability can be improved significantly. For indirect improvement, the two metrics of APAR and RPAR demonstrate the behavior of a group of ASs and their integrated impacts on the rest of the Internet. Combining these two metrics can further improve the overall routing scalability. The boundary values of these metrics can be used to evaluate how well the short-term goals have been achieved.

Moreover, various deployment strategies lead to different effects in achieving the short-term goals depending on the real demands. Multiple incentives decide the existence of various strategies and, hence, different short-term achievements.

2) On Long-Term Goals and Effects: The new findings of the evaluation also provide important guidelines for future solutions aiming at a long-term architecture evolution. For such considerations, the new architecture has to be open and extensible to accommodate demands and changes from ISPs and users, which may urge design principles different from the original ones [5], [28]. From a practical view, Internet should be able to evolve from the current through gradual steps.

Moreover, evaluating the current Internet is too difficult and inconvenient [18]. For long-term consideration, ***systematic monitoring and evaluation capabilities (methods and tools) should be provided and supported at the architecture level*** to identify problems and find needed changes. In other words, the AS-level analysis method is used in this paper as all other existing works do not have to worry about the limited data sources and the limited inference methods. Instead, the innovative method using quantitative metrics in this paper can be a norm and a part of the NIA itself.

D. Limitations of the Evaluation

Several limitations of our evaluation are worth discussing.

1) Our evaluation is based on some existing public BGP data sources. However, every BGP observation point has its limited visibility. Other useful data such as finer granularity traffic patterns of the ASs are hard to get because of the nondisclosure agreements of the Internet stakeholders [6].
2) The existing topology data and the prescreen algorithms we use may also not be accurate enough due to the extreme complexity of the interdomain routing system.
3) Internet is changing fast, and many latest trends may not be revealed in our work.
4) In deployment benefits' evaluation, we only consider some solutions, but future ones may deviate from them.

Due to these limitations, the evaluations are not 100% accurate. However, our goal is to present a reference and guide for the potential new architectures, which is the unique contribution of this paper as the first such effort.

VII. CONCLUSION

In this paper, we have tried to fill a gap between the designs of NIA and the evaluation efforts through an AS-level interdomain routing system evaluation. The major idea was systematically defining a series of quantitative metrics to reveal hidden information and observations that may be useful in improving the status and deploying candidate new architectural solutions. The results of the evaluation can be further applied to find the deployment strategies with balanced incentives for both customers and ISPs.

REFERENCES

[1] D. Meyer, L. Zhang, and K. Fall, "Report from IAB workshop on routing and addressing," RFC 4984, Sep. 2007.

[2] J. Pan, R. Jain, S. Paul, and C. So-In, "MILSA: A new evolutionary architecture for scalability, mobility, and multihoming in the future Internet," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1344–1362, Oct. 2010.

[3] J. Pan, S. Paul, and R. Jain, "A survey of the research on future Internet architectures," *IEEE Commun. Mag.*, vol. 49, no. 7, pp. 26–36, Jul. 2011.

[4] S. Paul, J. Pan, and R. Jain, "Architectures for the future networks and the next generation Internet: A survey," *Comput. Commun.*, vol. 34, no. 1, pp. 2–42, Jan. 2011.

[5] A. Feldmann, "Internet clean-slate design: What and why?" *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 3, pp. 59–64, Jul. 2007.

[6] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proc. ACM SIGCOMM*, Aug. 30/Sep. 3, 2010, pp. 75–86.

[7] Global Environment for Network Innovations (GENI) Project. [Online]. Available: http://www.geni.net/

[8] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," IETF, Fremont, CA, USA, IETF RFC 4271, Jan. 2006.

[9] L. Cittadini, W. Mühlbauer, S. Uhlig, R. Bush, P. François, and O. Maennel, "Evolution of Internet address space deaggregation: Myths and reality," *IEEE J. Sel. Areas Commun.*, vol. 28, no. 8, pp. 1238–1249, Oct. 2010.

[10] G. Huston, "Analyzing the Internet BGP routing table," *Internet Protocol J.*, vol. 4, no. 1, pp. 2–15, Mar. 2001.

[11] T. Bates, P. Smith, and G. Huston, The CIDR Report. [Online]. Available: http://www.cidr-report.org

[12] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 address allocation and the BGP routing table evolution," *ACM Comput. Commun. Rev.*, vol. 35, no. 1, pp. 71–80, Jan. 2005.

[13] IRTF Routing Research Group Charter. [Online]. Available: http://www.irtf.org/charter?gtype=rg&group=rrg

[14] K. Xu, Z. Duan, Z.-L. Zhang, and J. Chandrashekar, "On properties of Internet exchange points and their impact on AS topology and relationship," in *Networking*. Berlin, Germany: Springer-Verlag, 2004, pp. 284–295.

[15] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: Inference, modeling and generation," *IEEE Commun. Surveys Tuts.*, vol. 10, no. 2, pp. 48–69, 2nd Quart., 2008.

[16] G. Huston, "Scaling inter-domain routing—A view forward," *Internet Protocol J.*, vol. 4, no. 4, pp. 2–16, Dec. 2001.

[17] T. Bu, L. Gao, and D. Towsley, "On characterizing BGP routing table growth," in *Proc. IEEE Global Telecommun. Conf.*, Nov. 2002, pp. 2185–2189.

[18] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush, "10 lessons from 10 years of measuring and modeling the Internet's autonomous systems," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 9, pp. 1810–1821, Oct. 2011.

[19] IETF Host Identity Protocol (HIP) working group charter. [Online]. Available: http://datatracker.ietf.org/wg/hip/charter/

[20] IETF Locator/ID Separation Protocol (LISP) working group charter. [Online]. Available: https://datatracker.ietf.org/wg/lisp/charter/

[21] H. Ballani, P. Francis, T. Cao, and J. Wang, "ViAggre: Making Routers Last Longer!" in *Proc. Hotnets*, 2008, pp. 1–6.

[22] D. Farinacci, V. Fuller, D. Meyer, and D. Lewis, "Locator/ID separation protocol (LISP)," RFC 6830, Jan. 2013.

[23] NSF NeTS FIND Initiative. [Online]. Available: http://www.nets-find.net

[24] NSF Future Internet Architecture Project. [Online]. Available: http://www.nets-fia.net/

[25] CAIDA: The Cooperative Association for Internet Data Analysis. [Online]. Available: http://www.caida.org

[26] University of Oregon Route Views Project. [Online]. Available: http://www.routeviews.org

[27] S. Paul, R. Jain, J. Pan, and M. Bowman, "A vision of the next generation Internet: A policy oriented perspective," in *Proc. BCS Int. Conf. Vis. Comput. Sci.*, Sep. 22–24, 2008, pp. 1–14.

[28] C. Dovrolis, "What would Darwin think about clean-slate architectures?" *ACM SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 1, pp. 29–34, Jan. 2008.

[29] V. Cerf, B. Davie, A. Greenberg, S. Landau, and D. Sincoskie, FIND observer panel report, Apr. 9, 2009. [Online]. Available: http://www.nets-find.net/FIND_report_final.pdf

[30] Internet of Things. [Online]. Available: http://en.wikipedia.org/wiki/Internet_of_Things

[31] K. Fall, G. Iannaccone, and S. Ratnasamy, "Routing tables: Is smaller really much better?" in *Proc. HotNets*, Oct. 2009, pp. 1–6.

[32] P. Calhoun, T. Johansson, and C. Perkins, "Diameter Mobile IPv4 Application," RFC 4004, Aug. 2005.

[33] S. Kent and R. Atkinson, "Security architecture for the Internet protocol," RFC 2401, Nov. 1998.

Jianli Pan (S'08) received the B.E. degree from Nanjing University of Posts and Telecommunications, Nanjing, China, and the M.S. degree from Beijing University of Posts and Telecommunications, Beijing, China. He is currently a Ph.D. candidate in the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO, USA.

His current research includes: 1) next-generation Internet architecture and related issues such as routing scalability, mobility, multihoming, and Internet evolution; and 2) cyber-assisted energy efficiency, green building, smart energy, and sustainability.

Raj Jain (F'93) received the Ph.D. degree in applied math (computer science) from Harvard University, Cambridge, MA, in 1978.

He is currently a Professor of computer science and engineering with Washington University in St. Louis, St. Louis, MO, USA. Previously, he was one of the Cofounders of Nayna Networks, Inc., a next-generation telecommunication systems company in San Jose, CA, USA. He was a Senior Consulting Engineer with Digital Equipment Corporation, Littleton, MA, USA, and then a Professor of computer and information sciences with The Ohio State University, Columbus, OH, USA. He is the author of *Art of Computer Systems Performance Analysis*, which won the 1991 "Best-Advanced How-to Book, Systems" Award from the Computer Press Association, and of *High-Performance TCP/IP: Concepts, Issues, and Solutions* (Prentice Hall, November 2003), which is his fourth book. Recently, he has coedited *Quality of Service Architectures for Wireless Networks: Performance Metrics and Management* (April 2010).

Prof. Jain is a Fellow of the Association for Computing Machinery (ACM). He was a recipient of the ACM SIGCOMM Test of Time Award and the Center for Development of Advanced Computing-Advanced Computing and Communication Society Foundation Award in 2009. He ranks among the top 75 in CiteSeerX's list of Most Cited Authors in Computer Science.

Subharthi Paul (S'12) received the B.S. degree from the University of Delhi, New Delhi, India, and the Master's degree in software engineering from Jadavpur University, Kolkata, India. He is currently working toward the Doctoral degree in the Department of Computer Science and Engineering, Washington University in St. Louis, St. Louis, MO, USA.

His primary research interests are in the area of future Internet architectures, including software-defined networks, data center network architectures, and application delivery networking for cloud and multicloud environments.