Scalability Analysis of the TurfNet Internetworking Architecture

Jordi Pujol, Stefan Schmid, Lars Eggert and Marcus Brunner

Abstract — **TurfNet is an internetworking architecture that enables communication among autonomous and heterogeneous network domains. The architecture uses a global identity namespace and does not require global addressing or a common internetworking protocol. It integrates the new concept of dynamic network composition with other recent architectural concepts, such as decoupling locators from identifiers. This paper analyzes whether the TurfNet naming and interdomain routing mechanisms can scale to networks of the size of the global Internet. To the authors' knowledge, this is one of the first scalability analyses of any next-generation internetworking proposal based on today's Internet AS topology. It adapts existing research results on the topology of the Internet's autonomous system graph and related results that quantify typical traffic patterns to construct a realistic model of an Internet-like TurfNet.**

I. INTRODUCTION

THE basic principles of the original Internet architecture include end-to-end addressing, global routeability and a single address space of IP addresses that act as locators and node identifiers at the same time. These principles are suitable for static and well-managed network hierarchies. However, since the Internet has evolved from a small research network to a worldwide information exchange network, a growing diversity of commercial, social and governmental interests have led to increasingly conflicting requirements among the competing stakeholders. These conflicts create tensions that the original Internet architecture struggles to withstand. Clark *et al.* refer to this development as "tussles in cyberspace" [1]. These findings have prompted research into different internetworking architectures, such as FARA [2], Plutarch [3] or TRIAD [4].

Concurrently with this research into new internetworking architectures, a demand for private, autonomous networks is growing. One important aspect of this autonomy is address space control. Network Address Translation (NAT) [16] is a popular method for reusing address space and decoupling routing in the private network from routing in the public Internet. This enables network domains to change access providers or to multi-home by attaching to multiple service providers at the same time. NATs also hide the internal structure of private networks to the outside.

Braden *et al.* [5] propose the meta-architectural principle that individual regions of a network should be allowed to differ from each other in order to "minimize the degree of required global architectural consistency." The TurfNet architecture [8] (shortly summarized in Section II) adopts these principles as a necessary enabler for diversity between domains. It introduces a flat, unstructured host identity space that enables the use of different addressing and routing mechanisms in each individual autonomous network.

The use of a flat namespace raises important scalability and performance issues. This paper studies the performance of the TurfNet naming and interdomain routing mechanisms when they operate on realistic, large-scale topologies with one billion nodes. The internetwork topology and communication patterns are based on existing analyses of the Internet autonomous system (AS) graph and inter-AS communication patterns (Section III). The main contribution of this paper is the analysis of the scalability property (Section V) of the TurfNet naming and interdomain routing mechanisms, using the evaluation methodology explained in Section IV. These results can also provide an indication of the scalability property of other naming and interdomain routing proposals that use global identifiers for inter-domain routing. Finally, the paper concludes with a discussion of the analytical results (Section VI and VII).

II. THE TURFNET ARCHITECTURE

The TurfNet architecture focuses on enabling interoperation between otherwise autonomous networks. These autonomous networks are modularized according to the inherent boundaries drawn by the different interests of the stakeholders involved. This paper uses the name *turf* to denote such an autonomous network. The term turf has an innate connotation to ownership and responsibility that the TurfNet architecture reflects.

One key architectural feature of TurfNet is an explicit separation of node identifiers and node locators, similar to HIP [6], multi6 [7], or other proposals. TurfNet also uses an unstructured node identifier space that enables the use of different addressing and routing mechanisms in each autonomous turf. A mapping translates node identities into node locators suitable for network-layer data forwarding. The TurfNet architecture manages the global identifier spaces, whereas address spaces are local to each individual turf. This allows using different addressing and routing mechanisms in individual turfs.

Figure 1 illustrates the key components of TurfNet. The *turf control* is a logical, per-turf entity (centralized or distributed) that consists of a turf's essential control functions and services, such as address resolution and routing.

A *turf node* is a network node in a specific turf. For turf-local communication, the turf node must support the local network protocols and addressing schemes. A physical node can participate as a full-fledged turf node in multiple turfs at the same time, allowing multi-homing potentially using different

Jordi Pujol, Stefan Schmid, and Marcus Brunner are with the NEC Network Laboratories in Germany. Lars Eggert is with Nokia Research, the work has been performed while he was at NEC Network Laboratories.

technologies. Each turf node possesses one or more global identifiers, which map into turf-local locators used for addressing and routing within the local turf.

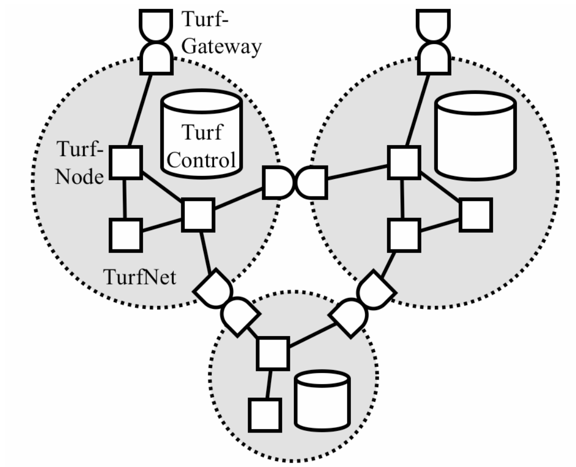

Figure 1. Key components of the TurfNet architecture.

Turf gateways are special, multi-homed turf nodes. Besides participating in multiple turfs at the same time, they can relay traffic between these different turfs. When turfs use different addressing or protocol mechanisms, the gateways also perform the required address and protocol translations when relaying traffic. For example, a gateway between IPv4 and IPv6 turfs translates between the two network protocols and their respective address spaces. Turf gateways enable interoperation by performing the necessary translation or emulation across independent turfs, if required.

TurfNet adopts a registration/lookup based scheme to enable communication based on node identifiers both within and across turf boundaries. A turf node becomes reachable by other nodes only after registering its identifier and intra-turf address with the local turf control. This registration propagates through the hierarchy of composed turfs to establish turf-external reachability. A turf control always forwards non-local registration messages to its parent turf controls (vertical). The result is a system that guarantees that lookups terminate at some point along the hierarchy to the top-level turf. Turfs may also choose to forward registrations across peering links (horizontal). This optional optimization can result in shorter lookup chains and reduces load on higher-level turfs at the expense of increased registration table sizes (see Section V).

When a turf node initiates communication, it attempts to look up the node identifier of the desired peer via the local turf control. As for registrations, turfs may optionally decide to also forward lookups across horizontal peering links in order to improve resolution probability, at the expense of an increase in lookup rates.

Figure 2 illustrates a registration and lookup operation in a scenario that involves peering interconnects; i.e., turfs not only forward registrations and lookups "up" the hierarchy but also "sideways" to peered turfs. Here, a registration for a node with identifier N propagates up the turf hierarchy (right side). Intermediate turfs register this node in their local turf control. Later, a lookup request "*N?*" for the node with identifier N appears (left side). This propagates upwards as well. The number of turfs a request is propagated along the peering interconnects is defined as *scope*, which is a parameter that can be chosen on a per requests basis by the client or by the local turf. Consequently, the lookup request that propagates up the left side can be resolved before it reaches the topmost level of the hierarchy, reducing load on those turfs.

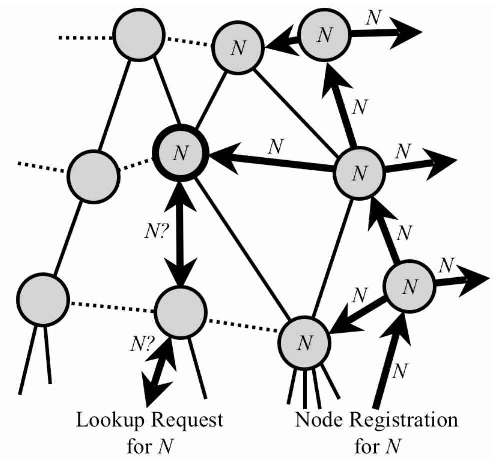

Figure 2. Registration and lookup.

Further details on the operation of TurfNet are provided in [19].

III. INTERNET CHARACTERISTICS

The TurfNet scalability analysis requires a model of a realistic, large-scale internetwork topology. This section describes the parameters of such a model, based on existing investigations of Internet topology, structure and communication characteristics.

To a certain degree, a turf is comparable to an autonomous system (AS) in the Internet. Both are administratively independent sub-domains interconnected into a global network through border gateways, but turf gateways may also perform locator or protocol translation. However, this paper assumes that the difference in operation does not influence the structure of the overall system significantly. It thus constructs a large-scale hierarchy of composed turfs that is similar in size and structure to the Internet AS-level topology by assuming that each Internet AS represents a separate turf.

Existing research results extract the AS-level topology of the Internet from the routing tables of the Border Gateway Protocol (BGP) [9][10]. Subramanian *et al.* [13] apply heuristics to the AS-level graph to determine inter-AS relationships (peering or customer/provider) between any two connected autonomous systems. Based on these heuristics and a recent BGP data set [13], the flat AS-level graph of the Internet can be transformed into a TurfNet hierarchy. This hierarchy has turfs at different *levels*, from 1 (the highest) to 5 (the lowest). All end systems are located at the bottom of the hierarchy at level 5. Turfs higher up the hierarchy do not have end systems and serve as providers of other turfs. At level 1 of the hierarchy, turfs have no providers and only customer turfs.

Table 1 illustrates some properties of the hierarchical graph obtained through this process. Turfs at level 1 form a connected graph that is almost a fully-meshed graph. Only a few level-1 turfs are not directly connected to all others, causing the diameter of the overall level-1 graph to be 2. At levels 2 and lower the peering graph of a given level becomes more disconnected. Within level 2, 193 out of the 215 turfs belong to the same highly connected graph the other ones belong to smaller strongly connected components. The mean path length between connected turfs is significantly longer (3.9 vs. 1.25 at level 1), indicating that the graph is not as densely connected as the level-1 one. The lower mean degree of 5.7 (vs. 15.8 at level 1) also illustrates this. On the other hand, level-2 turfs have many provider links to level-1 turfs; 10.66 per turf on average. Level 3 is even more sparsely connected with either none or a single peering link. Turfs at levels 3-5 have around 2-3 links to higher-level provider turfs.

Level	Turf Count	Mean Peering Path Length	Max. Diameter	Mean Degree per Turf	Mean Providers per Turf
1	22	1.25	2	15.8	0
2	215	3.90	10	5.7	10.66
3	1391	1.98	11	1.0	3.34
4	1421	(no peering)			2.31
5	13872				1.83

Table 1. Properties of an Internet-scale turf hierarchy.

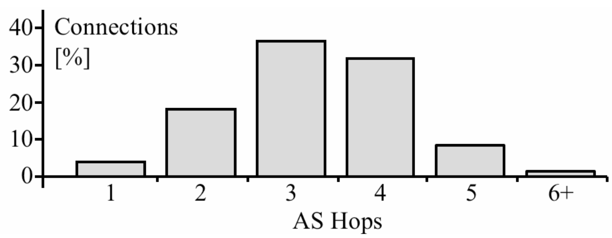

Figure 3. Internet communication distance distribution (Source: [11]).

To make realistic assumptions about the communication patterns of nodes in a large-scale TurfNet hierarchy, the analysis uses the results of previous studies of traffic patterns in the Internet [11][12]. They show that the typical end-to-end connection only traverses between 3-4 AS hops (see Figure 3). This paper assumes that a similar distribution is likely to be present in a comparable TurfNet topology.

IV. EVALUATION METHODOLOGY

The methodology for a scalability analysis of the TurfNet architecture defines the metrics, main assumptions and the mathematical model it is based on. By choice, a mathematical analysis was favored over simulations in order to evaluate the performance of a large-scale internetwork, which cannot be simulated with today's simulators.

Scalability of the TurfNet architecture is mainly limited by two factors. First, the size of the registration table, which contains information about the nodes inside each turf, in its customer turfs, and peering turfs. Second, the frequency of lookups for non-local nodes that the turf control must handle.

Especially for turfs located at higher levels in a turf hierarchy, the registration tables can grow rapidly, because they store information about all nodes in any customer turf at lower levels that would like to be globally reachable. Similarly, lookups for nodes generally propagate up the TurfNet hierarchy, and can cause high lookup frequencies at higher-level turfs, which again may render the architecture infeasible. Thus, the analysis focuses on these two metrics to investigate the scalability properties of TurfNet.

Although TurfNet enables nodes to selectively register with remote turfs to control from where they are reachable, the scalability analysis assumes that all nodes register for global reachability. This is the worst-case scenario for TurfNet given the above assumptions and hence the most interesting case for scalability, because it requires the top-level turfs to maintain state information for all nodes in the network.

Because the goal of the scalability analysis is to investigate TurfNet behavior for very large networks, it distributes 1 billion (10^9) nodes across the bottom, level-5 turfs of the hierarchy. This population exceeds current estimates on the size of the Internet [14]. For simplicity, the nodes are evenly distributed among level-5 turfs, *i.e.*, each of the 13,872 turfs at level-5 contains approximately 72,000 nodes. Note that an empirical analysis (not included here due to space restrictions) demonstrates that an exponential distribution of nodes to turfs does not significantly change the results.

Finally, the analysis assumes that nodes have a fixed, uniform communication initiation frequency of one communication instance to a different node (external to the local turf) every 100 seconds. The analysis assumes that this is a uniform constant for all nodes in all turfs, mainly because – to the authors' knowledge – no empirical data exists that allows making a more realistic choice.

V. SCALABILITY ANALYSIS

This section analyzes the scalability properties of the naming and routing mechanisms of the TurfNet architecture. The basis of this analysis is the large-scale TurfNet topology modeled after the Internet AS-level graph and the traffic characteristics described in Section III. The analysis is based on the assumptions defined in Section IV and carried out in *MatLab*. The metrics used in the analysis are the registration table sizes required at each turf to maintain node registrations and the rate of arriving lookup requests that a turf control has to handle.

To investigate the effects of different scopes, which determine how far registrations propagate across horizontal compositions, the analysis compares scenarios that use scopes of 0-2 hops. A scope of zero is the baseline case for TurfNet, in which peering links are never used for communication. With a scope of one, node registrations propagate only to immediate neighbors, and with a scope of two, they propagate also to neighbors up to two hops away. Note that the analysis uses the same scope at all levels and for all registrations, even though TurfNet supports the use of different scopes at different levels or for different registrations.

A second effect that this analysis investigates is the impact of different node lookup schemes. The first scheme investigates a scheme where turfs forward lookup requests up to all higher-level turfs. The scheme is compared to a variant where turfs only forward node lookups up ones, to their highest-level provider turf. This comparison quantifies the impact different lookup schemes can have on the overall lookup load.

Although the TurfNet hierarchy used for this analysis has five levels, with the top-level turfs at level 1 and all end hosts at level 5 at the bottom, the results below only discuss levels 1-4. This is, because no horizontal peering links exist between level-5 turfs and thus all inter-turf communication passes through turfs at level 4 and above. In addition, no inter-turf lookup requests arrive at level 5 or are resolved there.

Moreover, the results include a “level 0” turf that acts as the logical root of the TurfNet hierarchy. It is required for operation in the baseline case that uses a scope of zero, where peering links are not used.

The first part of this evaluation focuses on registration table sizes. Figure 4 compares the distributions of registration table sizes for turfs at different levels of the turf hierarchy (level 0-4) and for different scopes. The top graph shows the table sizes for a scope of zero, the middle graph shows them for a scope of one, and the bottom graph for a scope of two.

Figure 4 illustrates several important scalability aspects. First, it shows that an increasing scope eliminates the need for the level-0. Note that a level-0 turf may not be feasible due to policy reasons, as a single turf would constitute the “root” for the whole TurfNet. The level-1 registration tables hold between 100-500 million entries each with a scope of zero (top graph), *i.e.*, information about only 10-50% of the overall nodes. With a scope of one, they already hold between 820-920 million entries (middle graph), and with a scope of two, each level-1 turf contains information about all nodes in the system (bottom graph). In the latter case, they can thus replace the root turf by replicating its information.

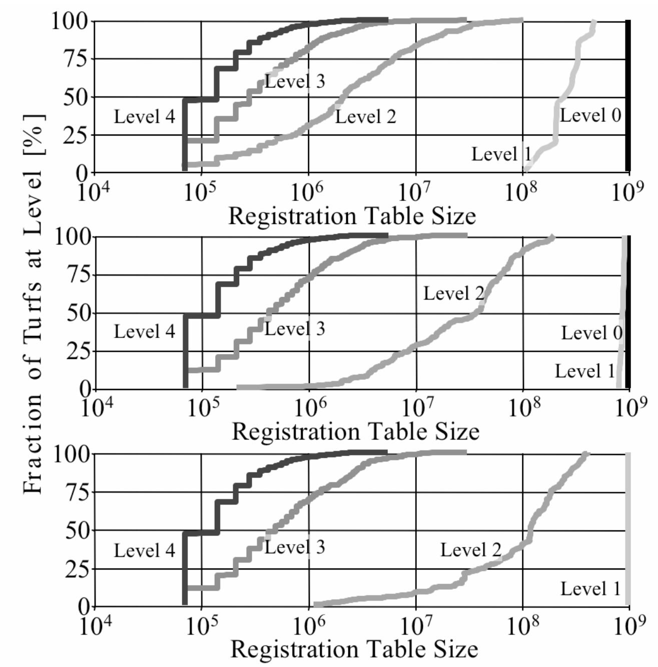

Figure 4. Distribution of registration table sizes for turfs at different levels with scopes of zero (top graph), one (middle graph), and two (bottom graph).

State tables at lower levels also increase. Median state tables at level-2 increase from 2.3 million entries with a scope of zero to 38 million with a scope of one and 123 million with a scope of two. However, in all cases the variance between tables at level-2 is much higher than at level-1. Level-1 tables have sizes that are within the same order of magnitude, whereas level-2 table sizes can vary over four orders of magnitude with a scope of zero. This illustrates that the node population below a level-2 turf can vary enormously. An increase in scope reduces this variance to two orders of magnitude with a scope of two, because node information from larger level-2 turfs is replicated at smaller ones, evening out the differences.

Unlike at levels 1 and 2, state tables at levels 3 and 4 do not significantly grow with an increasing scope. The median level-3 table has around 280 thousand entries with a scope of zero and 500 thousand with a scope of two. At level-4, the median state table has 14 thousand entries independent of the used scope. This result is due to the low occurrence of peering links at those levels, which render larger scopes ineffectual.

These results demonstrate that the use of peering links during node registration can significantly increase table sizes and consequently the resolution likelihood, especially at higher levels. The second part of this analysis investigates under which conditions larger scopes result in lookup benefits and under which they merely add overhead, *i.e.*, increase state tables without a corresponding reduction in lookup loads. The second part therefore focuses on the rate of arriving lookups.

Figure 5 shows the median arriving lookup request rate – with quartiles as error bars – across turfs at all levels of the hierarchy (different shading denotes scope). In the baseline case, with a scope of zero, approximately 23 million requests arrive at the level-0 root turf per second. With an increased scope of one, rates drop to around 140 thousand/seconds, because, as described above, now level-1 turfs already contain much of the overall registration information. A scope of two eliminates the need for the root level-0 turf, as all requests are resolved at level 1. (This however comes at the cost of a corresponding increase in registration table sizes.)

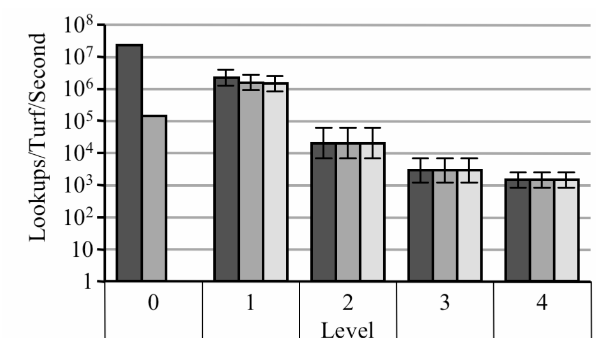

	0	1	2	3	4
Scope = 0	22,902,00	2,170,250	20,031	2,884	1,442
Scope = 1	142,160	1,551,300	20,028	2,884	1,442
Scope = 2		1,424,000	19,973	2,884	1,442

Figure 5. Median arriving lookup requests per turf (with quartiles) for each level of the hierarchy and different scopes.

At level 1, the gain in lookup rate reduction decreases. A scope of one reduces median lookup rates by approximately 30% from 2.1 million/sec to 1.5 million/sec; and a scope of two yields an additional, minor reduction to 1.4 million/sec. At other levels, the use of larger scopes does not significantly affect the median arriving lookup rate. The characteristics of the topology (**Table 1**) explain this difference. Whereas the average turf at level 1 has over 20 peering links, turfs at lower levels have much fewer peering links: 5-6 at level 2, approximately 1 at level 3, and none at levels below 3. Consequently, increasing the scope does not yield a benefit at levels 3 and 4.

These results identify one key issue: The drastic increase in registration table sizes caused by larger scopes does not translate into a similarly significant reduction in lookup rates. Figure 6 illustrates this finding by comparing the aggregate lookup rate arriving at each level, *i.e.*, the sum of the lookup rates arriving at all turfs belonging to a level. It shows that using a scope of one instead of zero only reduces the aggregate lookup load from 52 million lookups/sec to 40 million/sec, *i.e.*, a 20% decrease. Using a scope of two yields an even smaller decrease down to 37 million lookups/sec. Additionally, at levels 2-4, using a scope has no impact on the aggregate lookup load at all.

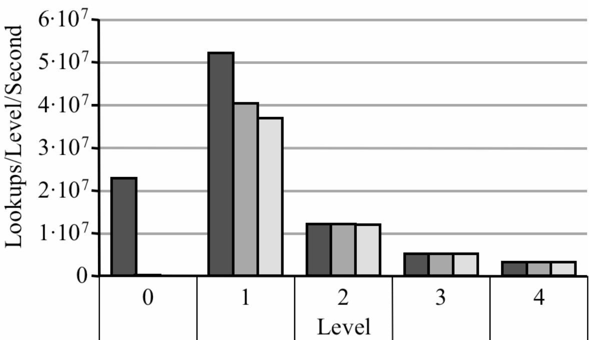

Level	0	1	2	3	4
Scope = 0	22,902,00	52,152,81	12,121,86	5,171,798	3,091,856
Scope = 1	142,160	40,405,43	12,107,39	5,171,798	3,091,856
Scope = 2	0	36,953,57	12,043,32	5,171,798	3,091,856

Figure 6. Aggregate arriving lookup requests for each level of the hierarchy and different scopes.

This result is surprising, because one would expect that due to the locality inherent in the communication patterns, lower levels resolve the majority of lookup requests. However, the simulation model assumes that a turf forwards any lookups it cannot resolve locally to all its higher-level provider turfs. This amplifies the lookup load at higher levels, even when individual turfs filter out duplicate requests, because even when a request is resolved along one lookup path, copies will still propagate up the hierarchy, causing significant amounts of load. Note, however, that this is a limitation in the initial model and not of TurfNet: the architecture only requires that one copy of a lookup request can reach the root to guarantee resolvability.

The remainder of this section evaluates an optimized lookup forwarding mechanism that significantly reduces lookup loads by exploiting the abovementioned redundancy. Compared to the “complete” variant discussed so far, the “greedy” registration variant only forwards node lookups up one vertical link to the highest-level provider turf, randomly choosing a single one if a turf has multiple providers of the same highest-level.

Figure 7 compares the distribution of lookup requests that arrive at a turf per second under the “complete” and “greedy” lookup schemes (top and bottom graph respectively),using a scope of two in both cases. At level 1, the greedy scheme results in a significant reduction of the lookup load: the median lookup rate drops from 1.4 million requests/second down to 195,000 requests/second, roughly an order of magnitude less. Furthermore, both the maximum and minimum request rate drop significantly, the minimum from 500,000 to 15,000 per second and the maximum from 3.5 million to 1.9 million per second. The improvement at level 2 is also significant: median lookup rates decrease from 20,000 to 5,700. Levels 3 and 4 see similar improvements.

These numbers illustrate that the avoidance of redundant lookups along vertical paths causes a significant reduction in lookup rates. Another obvious difference between the two schemes is that whereas all turfs see a minimum lookup load of approximately 500 requests/sec under the “complete” scheme, approximately 45% of level-4 turfs, 24% of level-3 turfs and 8% of level-2 turfs see no lookup requests at all under the “greedy” scheme. This is due to the multi-connected nature of the internetworking hierarchy, where many turfs have multiple connections to higher layers. Note, however, that the “complete” scheme does not necessarily result in a more balanced use of provider links either, because although it can return multiple path choices to the requesting node, it is up to the node to utilize them for data communication.

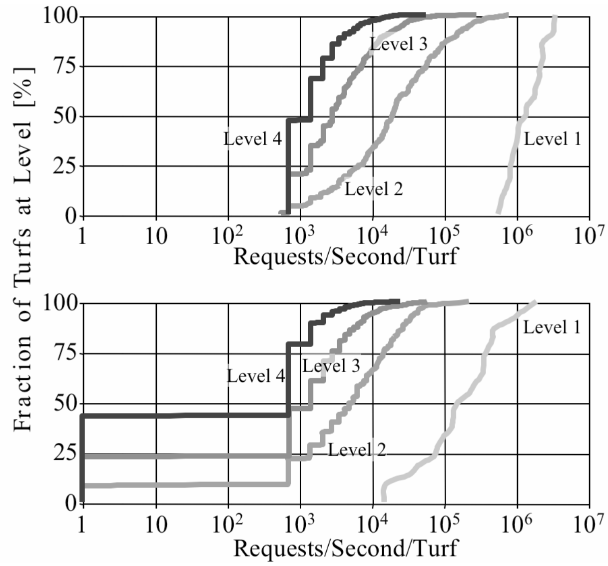

Figure 7. Distribution of arriving lookup requests per turf and second for different levels of the hierarchy and a scope of 2, under the “complete” lookup scheme (top) and the “greedy” lookup scheme (bottom).

VI. DISCUSSION

The primary constraints of the TurfNet architecture are the storage space requirements of node registrations as well as the registration and lookup request loads at the turf controls. Both of these affect the higher-level turfs in a hierarchy most critically, because those turfs need to maintain registrations and possibly serve lookup requests for large fractions of the overall node population. The use of peering interconnects between turfs at the same level, *i.e.*, to forward registrations or lookups, is effective in reducing lookup loads at higher-level turfs. However, their use increases registration table sizes significantly and does not cause a similarly significant reduction in lookup rates. For example, peering only reduces the median lookup rate at level-1 turfs by approximately 34% when increasing the scope from zero to two (see Figure 5).

Although the gains of increased scopes are not dramatic, they have two important advantages. First, they allow nodes to find shorter inter-turf, end-to-end routing paths. This can decrease subsequent data traffic loads on the high-level turfs. Second, using a scope of two for lookups at level 1 of the hierarchy has the advantage that it eliminates the need for the root level-0 turf.

With “greedy” lookups with a scope of 2, the results in Section V show that turfs at level 1 need to store registration state for the entire population of the internetwork. They also need to answer a median of 195 thousand lookups/second, up to a maximum of 1.9 million lookups/second (see Figure 5). The majority of turfs at level 2 also maintains a large number of registrations and sees a high load of lookup requests.

It is important to note that lookup rates pose a more serious scalability issue than registration state sizes. For example, if a registration record, which consists primarily of a node identifier, local address plus some timestamps for soft-state and cache maintenance, demands 50 bytes of storage, the storage requirement for the complete level-0 registration table is 50 GB. These storage requirements are well within the capabilities of current distributed database systems. Consequently, turf control scalability at higher levels mostly depends on the arriving lookup loads (see Figure 7). The use of distributed system principles allows turf controls at those levels to cope with the required lookup loads. For example, existing database clusters that simply replicate the state information would be able to resolve the lookup requests at the estimated rates. Considering that today's commercial DNS servers can already handle tens of thousands of DNS queries per second, a cluster of 10-100 replicas, depending on arriving lookup rates, could handle the load of a level-1 turf. Alternatively, Distributed Hash Table (DHT) principles [15] could be exploited to distribute the storage and lookup load across the required number of servers which at the same time allows for low cost management and easy scale up. For example, a DHT consisting only of the required number of servers to handle the lookup load, and full routing knowledge at any node of the DHT would allow for very efficient load distribution.

Furthermore, the analysis in Section V also assumed that all nodes are equally popular. However, in the current Internet, a very small fraction of the overall node population – the "servers" – is the target of the vast majority of communication instances. The large volume of communications to these nodes can greatly benefit from cached lookup results of prior communications, which can significantly reduce lookup loads.

Finally, TurfNet lookup performance can be further optimized through speculative, push-down and caching of the registration state of popular static nodes at lower levels of the hierarchy, similar to techniques proposed for web caching [18]. This can reduce lookup delays for the most common communication destinations.

VII. CONCLUSION

TurfNet is a novel internetworking architecture that enables communication among highly autonomous and heterogeneous network domains. The architecture uses a flat, unstructured global identifier space and does not rely on global addressing. This paper examined whether the TurfNet architecture can scale to networks of similar structure and size than the global Internet. The scalability analysis suggests that the TurfNet architecture is generally feasible using current technology; core turfs located at higher levels of the hierarchy require replicated turf controls to sustain the arriving lookup loads.

Future work focuses on extending the scalability analysis to interdomain topologies that have different structures from the Internet-like one investigated in this paper. Additionally, different communication patterns, turf population distributions, and caching effects will be evaluated. One interesting question is whether the current static classification of links into peering and customer/provider types can be relaxed. TurfNet may be able to auto-tune its topology by automatically changing the classification of links based on observed communication patterns or other information. Finally, next steps include experimentation with a prototypical TurfNet system to evaluate the performance at packet processing and protocol levels.

REFERENCES

[1] D. Clark, J. Wroclawski, K. R. Sollins and R. Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. Proc. *ACM SIGCOMM*, Pittsburgh, PA, USA, August 2002, pp. 347-356.

[2] D. Clark, R. Braden, A. Falk and V. Pingali. FARA: Reorganizing the Addressing Architecture. Proc. *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, Karlsruhe, Germany, August 2003, pp. 313-321.

[3] J. Crowcroft, S. Hand, R. Mortier, T. Roscoe and A. Warfield. Plutarch: An Argument for Network Pluralism. Proc. *ACM SIGCOMM Workshop on Future Directions in Network Architecture (FDNA)*, Karlsruhe, Germany, August 2003, pp. 258-266.

[4] D. R. Cheriton and M. Gritter. TRIAD: A Scalable Deployable NAT-based Internet Architecture. Stanford Computer Science Technical Report, 2000.

[5] R. Braden, D. Clark, S. Shenker and J. Wroclawski. Developing A Next-Generation Internet Architecture. *Whitepaper* (http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.ps), July 2000.

[6] R. Moskowitz and P. Nikander. Host Identity Protocol Architecture. *Work in Progress* (draft-ietf-hip-arch-03), August 2005.

[7] J. Abley, B. Black and V. Gill. Goals for IPv6 Site-Multihoming Architectures. *RFC 3582*, August 2003.

[8] S. Schmid, L. Eggert, M. Brunner and J. Quittek. Towards Autonomous Network Domains. Proc. 8th *IEEE Global Internet Symposium*, Miami, FL, USA, March 17-18, 2005, pp. 13-18.

[9] H. Chang, R. Govindan, S. Jamin, S. Shenker and W. Willinger. On Inferring AS-level Connectivity from BGP Routing Tables. Proc. *IEEE INFOCOM*, New York, NY, USA, June 23-27, 2002.

[10] Z. Ge, D.R. Figueiredo, S. Jaiswal and L. Gao. On the Hierarchical Structure of the Logical Internet Graph. Proc. *SPIE ITCOM*, Denver, CA, USA, August 2001.

[11] S. Uhlig and O. Bonaventure. Implications of Interdomain Traffic Characteristics on Traffic Engineering. *European Transactions on Telecommunications*, Special Issue on Traffic Engineering, 2002.

[12] D. Krioukov, K. Fall and X. Yang. Compact Routing on Internet-Like Graphs. Proc. *IEEE INFOCOM*, Hong Kong, China, March 7-11, 2004.

[13] L. Subramanian, S. Agarwal, J. Rexford and R.H. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. Proc. *IEEE INFOCOM*, New York, NY, USA, June 2002.

[14] ISC. Internet Domain Survey. http://www.isc.org/ds/

[15] I. Stoica, R. Morris, D. Karger, M. Frans Kaashoek and H. Balakrishnan. Chord: A Scalable Peer-to-Peer Lookup Service for Internet Applications. *IEEE/ACM Transactions on Networking*, Vol. 11, No. 1, February 2003, pp. 17-32.

[16] P. Srisuresh and M. Holdrege. IP Network Address Translator (NAT) Terminology and Considerations. *RFC 2663*, August 1999.

[17] M. Caesar, T. Condie, J. Kannan, K. Lakshminarayanan, I. Stoica, S. Shenker, ROFL: Routing on Flat Labels, To appear: Proc. *ACM SIGCOMM*, Pisa, Italy, 2006.

[18] J. Touch and A. S. Hughes. The LSAM Proxy Cache - a Multicast Distributed Virtual Cache. *Computer Networks and ISDN Systems*, Vol. 30, No. 22-23, November 1998, pp. 2245-2252

[19] S. Schmid, L. Eggert, M. Brunner and J. Quittek. TurfNet: An Architecture for Dynamically Composable Networks. Proc. 1st IFIP International Workshop on Autonomic Communication (WAC), Germany, 2004.