

ParaNets: A Parallel Network Architecture for Challenged Networks

Khaled A. Harras, Mike P. Wittie, Kevin C. Almeroth, and Elizabeth M. Belding
 Department of Computer Science, University of California, Santa Barbara, CA 93106
 {kharras, mwittie, almeroth, ebelding}@cs.ucsb.edu

Abstract

Networks characterized by challenges, such as intermittent connectivity, network heterogeneity, and large delays, are called “challenged networks”. We propose a novel network architecture for challenged networks dubbed *Parallel Networks*, or, *ParaNets*. The vision behind ParaNets is to have challenged network protocols operate over multiple heterogeneous networks, simultaneously available, through one or more devices. We present the ParaNets architecture and discuss its short-term challenges and long-term implications. We also argue, based on current research trends and the ParaNets architecture, for the evolution of the conventional protocol stack to a more flexible cross-layered *protocol tree*. To demonstrate the potential impact of ParaNets, we use Delay Tolerant Mobile Networks (DTMNs) as a representative challenged network over which we evaluate ParaNets. Our ultimate goal in this paper is to open the way for further work in challenged networks using ParaNets as the underlying architecture.

I. INTRODUCTION

The increase in user demand and mobility, along with the evolution of wireless devices, have led to a new set of networking challenges. Some of these challenges include network heterogeneity, intermittent connectivity, large delays, the high cost of infrastructure deployment, and the absence of an end-to-end path. Network environments characterized by one or more of these challenges are dubbed “challenged networks” [1]. Examples of challenged network environments include battlefields, disaster relief efforts, and remote disconnected villages.

To address these challenges, new research thrusts in various challenged networks have emerged [1], [2], [3], [4], [5]. The DTN architecture, for example, works well by addressing common problems among different networks in challenged environments [1]. Conversely, other specialized solutions are tailored to the problem they try to solve, and therefore, are more customized to their designated environments. However, all of these solutions are based on the idea of operating over a single homogenous network, or sequential heterogeneous overlay networks.

In parallel to the increase in challenged network environments, we have begun to witness the inevitable convergence of different networking technologies. This convergence occurs by providing communication alternatives to users through carrying multiple devices, or a single device, with access to multiple networks [6], [5]. However, network protocols that operate on these devices are primarily designed to operate over a single network at a given time, or multiple channels within the same network. Most, if not all, current approaches for providing interoperability between heterogeneous networks rely on a high level overlay protocol that performs protocol translation. These overlays, however, are usually at network gateways rather than endpoints. We believe we can exploit the current and expected future convergence of networking technologies to better serve challenged networks.

Based on our vision, as well as the current trends in challenged networks research and networking technologies, we propose the *Parallel Networks (ParaNets)* architecture. The idea behind parallel networks is to provide an architecture over which network protocols, developed for challenged networks, can seamlessly utilize multiple heterogeneous networks in parallel. Each network can then be used as a channel for the protocol being used. Message types that are best suited for a given network are seamlessly sent using the appropriate channel.

The ParaNets architecture has several short-term research challenges and long-term implications which we discuss in this paper. Some of the short-term challenges include revisiting the transport, routing, addressing, security, and administrative issues in challenged networks. The major long-term implication of ParaNets is that the conventional protocol stack evolves into a more flexible and adaptive *cross-layered protocol tree*. We believe that merging this new concept with the current trend of cross-layer approaches [7], opens the door for protocols capable of providing more robust, timely, and intelligent decisions for challenged networks.

In order to quantitatively show the gain of using ParaNets, we study its impact on Delay Tolerant Mobile Networks (DTMNs) [4]. We use DTMNs as a representative challenged network to evaluate the ParaNets architecture. We validate our short-term challenges and long-term implications by showing the significant improvement in the results of the ParaNets-based solutions when compared to current approaches. Our main goal in this paper is to provide a novel architecture over which future challenged networking protocols can be built.

II. RELATED WORK

While early research attempts have addressed some problems in challenged networks, the term “challenged networks” was first introduced by Fall when he proposed a Delay Tolerant Networking (DTN) architecture for heterogeneous networks in extreme environments [1]. Other work on more specific issues, such as using “ferries” for message delivery [2], limiting message floods in sparse mobile environments [4], or tracking wild life patterns [3], have also been introduced. Most of this work, however, focuses on the particular challenged network being studied, in order to deliver and implement a particular solution efficiently designed for a specific domain.

The ParaNets architecture distinguishes itself from previous work, particularly the DTN architecture [1] in several ways. We exploit the current and expected future availability of multiple networks at the endpoints to improve the performance of current solutions. We also highlight and discuss the short-term challenges and long-term implications of ParaNets. Last, we take the first steps towards providing insight as to how these challenges can be addressed.

III. SYSTEM ARCHITECTURE

We discuss in this section two major issues. We first present the ParaNets architecture and show how endpoints can take advantage of multiple heterogeneous networks in parallel. Afterwards, we define a set of short-term challenges and long term implications that arise as a result of ParaNets.

A. The Parallel Networks Architecture (ParaNets)

We propose the ParaNets architecture based on the current and future expected trend of individuals accessing multiple networks through one or more devices. We exploit the availability of these devices through their participation in the creation of protocols that can better serve challenged networks. This idea has a profound impact on the services and protocols required to run on such devices and networks.

An instance of the ParaNets architecture is shown in Figure 1. The endpoints, A and B , have access to three different networks in parallel. The three networks are a classical DTN network [1], a cellular network, and a satellite network. The DTN network represents the challenged network in this case. This network has multiple regions with DTN gateways at the edges of these regions. These gateways perform translation and persistent store-and-forward relaying of “messages bundles” [1]. The DTN network is generally characterized by high bandwidth, but intermittent, and non-end-to-end connectivity. The cellular network, however, is characterized by continuous, end-to-end connectivity, but low bandwidth connections. Finally, the satellite network generally falls in the middle in terms of available bandwidth. It provides intermittent, end-to-end connectivity.

When nodes A and B need to transmit bundle messages over the DTN network, these messages are typically transmitted over the classical DTN network using store-and-forward relaying. Without ParaNets, functionality, such as, routing or reliability, are only provided via the challenged network. Given the characteristics of the challenged network, it becomes increasingly hard to develop robust, efficient solutions. With ParaNets, however, the available parallel cellular and satellite networks are treated as channels that can be used depending on the type of message. For example, the cellular network is best used for transmitting lightweight control information such as acknowledgements, synchronization messages, or routing updates. The satellite network, which usually provides predictable, scheduled, end-to-end connections, can also be used for transmitting control information, or small to average sized message bundles. We finally note that access to the parallel networks is not restricted to the endpoints, but can be extended to the DTN gateways as well, thus enabling extra services for communication in such environments.

B. Research Challenges and Implications

There are numerous new research challenges and implications that emerge as a result of ParaNets. Many of these challenges are open for future research. In this section, we identify some of these challenges. We also provide some insight on possible approaches to solve some of these problems. We generally classify these challenges into short-term challenges and long-term implications.

1) *Short-term challenges:* We now discuss what we believe are the most immediate and crucial research challenges created as a consequence of ParaNets. These challenges include:

Transport: The way in which various transport layer services are currently provided in challenged networks will most likely change with the introduction of ParaNets. Control messages related to these transport services, such as connection establishment, congestion control, and reliability, no longer need to traverse the challenged network. These messages now have an alternative channel over which they can be transmitted. Lightweight data

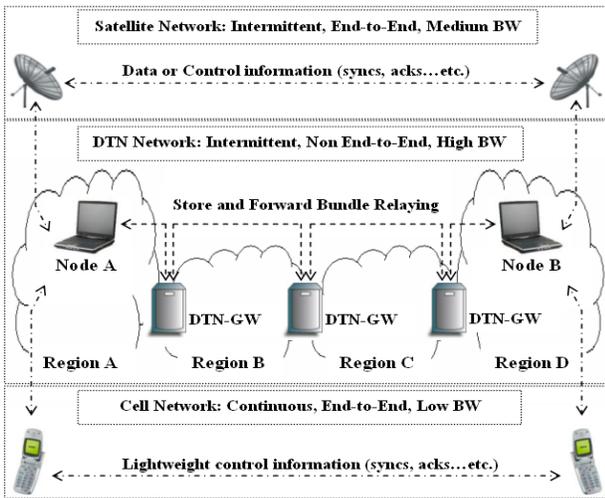


Fig 1. A general parallel network architecture.

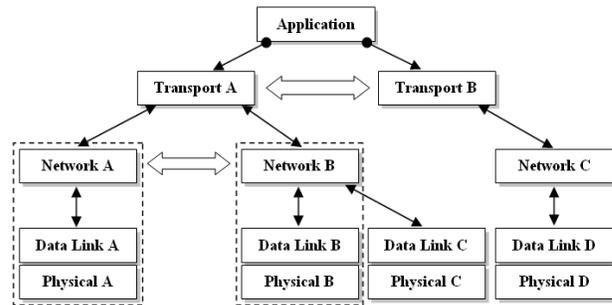


Fig 2. The new vision of a protocol tree.

carrying information regarding buffer sizes, acknowledgments, lost packets, and battery power, for example, can now be quickly delivered to far ends of a challenged network. We examine this approach for maintaining end-to-end reliability in our evaluation section.

Routing: Most of the routing approaches adopted in challenged networks rely on opportunistic or scheduled links that are available for short periods of time. Routing decisions are made based on information that can easily be stale due to the nature of the challenged network itself. Information is often disrupted or delayed for long periods of time. ParaNets, in this case, offer much more up-to-date information regarding the challenged network. This information can then lead to more efficient and reliable routing decisions. For example, with ParaNets, information such as node power, path buffers, routes, or location coordinates, can all be sent more quickly and reliably.

Addressing: Since nodes are connected to multiple networks, we need to develop methods by which to correlate and identify the different network addresses given to a node. In other words, we need to ensure that a given node, X , with a given IP, α , is also reachable through the cell number, β , and γ on a third network. New naming schemes and address resolution techniques that can map α to β to γ is yet another interesting research challenge.

Security: Since ParaNets allow operation over multiple networks, there are multiple paths over which security related attacks can be launched. However, these multiple paths also offer alternatives for ensuring authenticity and integrity. Security related information, such as certificates or various keys, can now be transmitted over an appropriate alternate path. Coordinating security measures across such heterogeneous environments is the important challenge.

Administration: Because of ParaNets, a node has access to multiple networks that fall under different administrative systems and domains. Each of these domains has its own operational procedures, administrative policies, and costs. For example, the traversal of traffic on a network, A , in order to provide services for a network, B , might not be allowed by A . The problem becomes more challenging if A is not designed for such scenarios. For example, the use of the cellular network to send control information might be rejected by cell providers offering voice-only service. The architecture we present takes the current trend of network convergence to new levels with new administrative challenges.

2) *Long-term implications:* Most protocols are guided and limited by the conventional protocol stack. While the idea of the vertically layered protocol stack has its advantages, this rigid architecture has been broken on many occasions to improve performance. This deviation is particularly prominent in recent research thrusts in cross-layer solutions [7], challenged overlay networks [1], and some multi-network devices [6]. We take these current trends, combine them with ParaNets, and present a vision of how we believe future protocols can be built.

We believe that future protocols will be built based on a tree, shown in Figure 2, rather than a fixed stack. The general idea behind the protocol tree is to build protocols in a more flexible and adaptable manner. The layers used in the protocol stack are still the same, but the ways in which they interact are different. Each layer in the tree is built such that the interface between layers are more adaptable to a range of possible underlying layers. When a session needs to be established, the most appropriate path, from root to leaf, is chosen, which then represents a temporary virtual stack for the session.

An application is able to use protocols provided by transport layers A or B , as shown in Figure 2. In this case, transport protocols A and B are best suited for different underlying networks. Furthermore, each transport protocol has the option of using multiple underlying networks. This is illustrated in Figure 2 where transport A can access both networks A and B . We believe that the network, MAC, and physical layers are more tightly coupled in such an architecture. This flexibility in the protocol tree leads us to develop more expandable and backwards compatible protocols, and also provides multiple opportunities for each layer to perform its tasks in the most efficient and effective manner.

The arrows shown in Figure 2 represent two concepts. The thin black arrows show how the current trend of cross layering will be maintained in the protocol tree. Sharing information vertically between layers has proven to be very useful in many cases [7], and we believe that this trend will continue, particularly for mobile and challenged networks. On the other hand, the thick horizontal arrows represent what we call cross-network, or cross-stack, cooperation. For example, if a device has connectivity to both networks A and B , routing information gathered from each network can be shared, enabling both layers to learn more than it is capable of knowing on its own. This fact becomes more apparent when we consider ParaNets combined with challenged networks. A node can learn much more about a route to other nodes in the challenged network through the parallel cellular network. We note that the protocol tree vision is more of a long-term implication of both, current research trends, as well as ParaNets. However, more attention needs to be directed towards this implication since we believe that it is a highly probable evolution of the current state of network design.

IV. EVALUATION

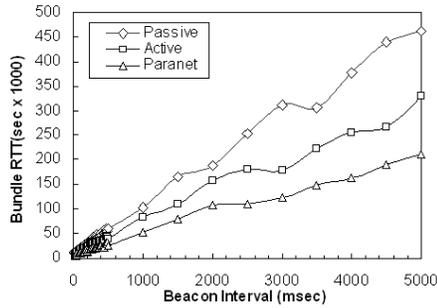
The primary goal of our evaluation is to demonstrate the impact ParaNets can have on challenged networks. To fulfill this goal, we use Delay Tolerant Mobile Networks (DTMNs) as the representative challenged network over which we conduct our evaluation. DTMNs are a special kind of DTNs, where all nodes are mobile, and no end-to-end path necessarily exists between any two nodes [4]. In our evaluation, we incorporate and address some of the short-term challenges mentioned in the previous section, such as reliability and message delivery.

A. Simulation Environment

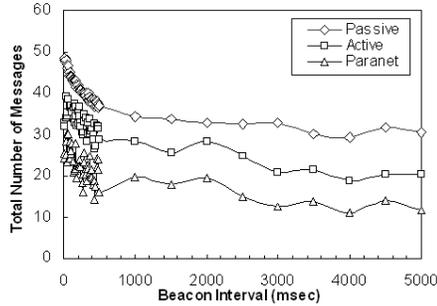
We conducted our evaluation using OPNET. We integrated the ParaNets architecture with that of DTMNs by adding the ParaNets handler. This handler coordinates access to multiple networks, which, for our evaluation, were an 802.11-enabled DTMN and a cellular network. We evaluated the performance of the *best* approaches used for message delivery and reliability in a classical DTMN, shown in previous research [4] [8], to that of the ParaNet-enabled DTMN. The best approaches we considered were the *passive*[4] and *active*[8] cures. These cures are approaches used to control and stop message floods in DTMNs by gradually healing message-infected nodes, either through a “passive” kill message or an “active” receipt. The passive and active cures are also used for end-to-end reliability in DTMNs. We particularly show how a ParaNet-enabled DTMN helps spread either one of these cures in a more efficient manner.

In our simulations, we used the random way-point mobility model, but avoided the major problem of node slow down. We used random way-point because our investigation showed that, for this particular case, the mobility model did not affect the relative performance of our solutions. The node speed ranged between 15 to 25 meters per second, and the rest period was 10 seconds. Every point in our results was taken as an average of 20 different seeds. We ran the simulation until the acknowledgment for message delivery was received.

We summarize the major parameters we use in our simulations as follows: *Terrain* is the area over which the *Number of Nodes* are scattered. We run our tests over $5km^2$ with nodes between 25 and 75 (nominal value = 50 nodes). The *Beacon Interval* is the period after which a beacon, used for neighbor discovery, is sent. The beacon interval range is between 10 milliseconds and 10 seconds (nominal value = 0.5 seconds). The *Times-To-Send* (TTS) is the number of times a node will successfully forward a message before it is dropped. This value ranges from 1 to 75 times (nominal value = 50). *Retransmission Wait Time* represents the amount of time a node remains idle after successfully forwarding a message to another node. When the retransmission wait time expires, the node then tries to resend the same message. This time ranges between 0 to 50 seconds (nominal value = 5 seconds). Finally, the *ParaNets Percentage* is the percentage of nodes that have access to a parallel network in order to make use of the ParaNets architecture. This percentage ranges between 2% and 100% (nominal value = 100%).

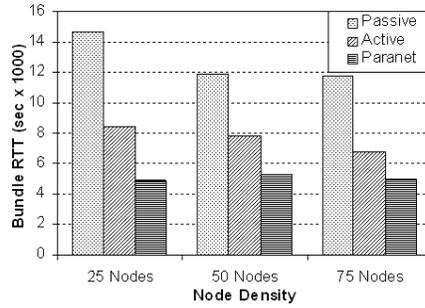


(a) Delay

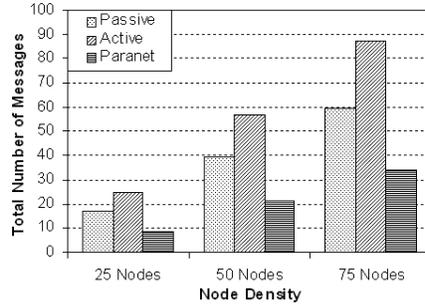


(b) Cost

Fig 3. Impact of beacon interval

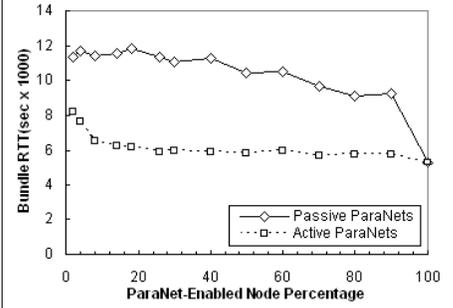


(a) Delay

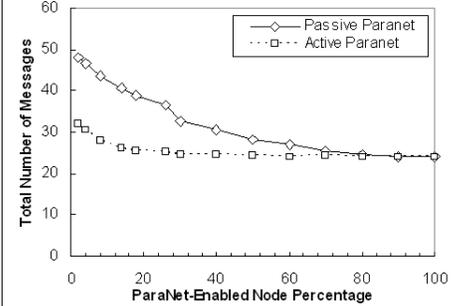


(b) Cost

Fig 4. Impact of number of nodes



(a) Delay



(b) Cost

Fig 5. Impact of ParaNets percentage

We use two metrics to evaluate the impact of ParaNets on DTMNs. The first metric is *Delay*, measured by the total time spent to send a bundle message as well as receive its acknowledgement through the cure. We also call this the round trip time of a message. The second metric is *Cost*, measured by the total number of bundle messages introduced into the network.

B. Results

We now present a summary of our extended set of simulations, along with a specific subset of our simulation results that clarify and support our goal. All of the results are shown for a single sender sending one bundle message to a single destination. The main goal of these experiments is to study the impact of ParaNets on DTMNs, which helps us better understand its potential impact on challenged networks.

Since both the beacon interval and times-to-send represent the willingness of a node (i.e., how hard a node tries to send a message), we only show results for the beacon interval. The results in Figure 3 show the impact of changing the beacon interval. We witness similar performance trends in the three approaches in terms of delay (Figure 3(a)) and cost (Figure 3(b)). In general, we incur an overall increase in round trip time as the beacon interval increases, and a decrease in the total number of bundles transmitted. The less frequently nodes send beacons, the harder it is for them to discover neighbors to which they can forward a message, resulting in an increase in delay. This loss of connection opportunities, however, is contrasted by an overall decrease in cost since messages are relayed less often.

The second result set is illustrated in Figure 4. The figure shows the impact of varying the network density (25, 50, and 75 nodes) on our metrics. For all approaches, as the number of nodes increases, the round trip time decreases and the total number of bundles transmitted increases. As the density increases, more nodes in the network spread the cure faster. This faster spread of the cure, however, comes at the cost of an increase in the total number of messages injected into the network.

In the results in Figures 3 and 4, we see that the ParaNets-based approach outperforms both the active and passive cure approaches regardless of the change in beacon interval or node density. The reason for this improvement in performance when using ParaNets is that the cellular network acts as a high-speed and low-bandwidth channel. This channel enables the cure to propagate much faster, which leads to two crucial results. First, the round trip time is greatly reduced since the return time is essentially zero because the acknowledgement/cure traverses via the cellular network instead of the challenged network. Second, the total number of messages is much smaller with ParaNets because, once the destination receives the message, the protocol using the ParaNets architecture quickly

stops nodes in the network from forwarding the message any further. In other words, the nodes in the network know earlier, via the parallel network, that the message was delivered, and therefore, stop retransmitting the message.

We note that the active approach outperforms the passive approach for both metrics in Figure 3. When the cure is “actively” transmitted as a separate message, it reaches the source faster since the uninfected nodes also participate in delivering the active cure. This reduction in return time also stops the nodes in the network from retransmitting earlier than in the passive approach, resulting in a smaller total number of messages transmitted. The less chatty passive approach, however, is more advantageous than the active one if we include the acknowledgement message as part of the cost, as shown in Figure 4(b).

The results we have presented so far assume that all the nodes in the network are ParaNets-enabled. This assumption also explains why we show only one line for ParaNets, because the active and passive approaches running over ParaNets perform the same; whether the cure is active or passive, it reaches all nodes as soon as the message is delivered to the destination. What happens, however, if only some nodes in a challenged network have access to a parallel network?

Figure 5 answers this question by showing the impact of having only a subset of nodes ParaNets-enabled. We ensure that the sender node is *not* ParaNets-enabled or else there will be no change in RTT since the sender will receive the cure as soon as the message is delivered to the destination. We now see a difference in the active vs. passive performance over the ParaNets architecture. The convergence in performance that we see is, again, because as more nodes are ParaNets-enabled, the faster the source receives the cure until no difference exists when all nodes are ParaNets-enabled. Overall, there is a decrease in both delay (Figure 5(a)) and cost (Figure 5(b)) as the percentage of ParaNets-enabled nodes increases. The interesting observation, however, is that similar improvement in performance is also achieved only with a small number of ParaNets-enabled nodes. This result occurs, as shown, for either the active or passive cure running over the ParaNets architecture.

V. CONCLUSIONS AND FUTURE WORK

The ParaNets architecture aims to provide a basis on which solutions targeted towards challenged networks can be built. Its design exploits the availability of multiple networks in parallel in such a way to help develop more efficient and robust solutions for challenged networks. We have studied the impact of the ParaNets architecture on a representative challenged network, namely, Delay Tolerant Mobile Networks (DTMNs). Solutions based on ParaNets are shown, through our evaluation, to outperform current state of the art solutions, even when only a subset of nodes are ParaNets-enabled.

This work represents the first steps towards developing a large scale, robust, and scalable architecture for challenged networks. This architecture must be flexible and expandable enough to be the basis upon which future solutions for challenged networks can be built. Many other problems and challenges need to be addressed in ParaNets. We have contributed by indicating some short-term research challenges, such as, transport, routing, addressing, security, and administration. We have also demonstrated the evolution of the protocol stack to a cross-layered protocol tree, as a long-term implication of ParaNets. Many challenges, however, remain to be tackled in future work. Nevertheless, we believe that most of the current challenged networks solutions can benefit from the incorporation of ParaNets into their architecture.

REFERENCES

- [1] K. Fall, “A Delay-Tolerant Network Architecture for Challenged Internets,” in *ACM SIGCOMM*, Karlsruhe, Germany, August 2003.
- [2] W. Zhao and M. Ammar and E. Zegura, “Controlling the Mobility of Multiple Data Transport Ferries in a Delay-Tolerant Network,” in *IEEE INFOCOM*, Miami, FL, March 2005.
- [3] P. Juang, et al., “Energy-Efficient Computing for Wildlife Tracking: Design Tradeoffs and Early Experiences With ZebraNet,” in *International Conference on Architectural Support for Programming Languages and Operating Systems*, San Jose, CA, October 2002.
- [4] K. Harras, K. Almeroth, and E. Belding-Royer, “Delay Tolerant Mobile Networks (DTMNs): Controlled Flooding Schemes in Sparse Mobile Networks,” in *IFIP Networking*, Waterloo, Canada, May 2005.
- [5] D. G. Andersen, H. Balakrishnan, and F. Kaashoek, “Improving Web Availability for Clients with MONET,” in *2nd Symp. on Networked Systems Design and Implementation (NSDI)*, Boston, MA, May 2005.
- [6] R. Chandra, P. Bahl, and P. Bahl, “MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card,” in *IEEE INFOCOM*, Hong Kong, March 2004.
- [7] V. Srivastava and M. Motani, “Cross-layer design: a survey and the road ahead,” *IEEE Communications*, vol. 43, no. 12, pp. 112–119, 2005.
- [8] K. Harras and K. Almeroth, “Transport Layer Issues in Delay Tolerant Mobile Networks,” in *IFIP Networking*, Coimbra, Portugal, May 2006.