Deafness and Virtual Carrier Sensing with Directional Antennas in 802.11 Networks

Michael Neufeld and Dirk Grunwald
University of Colorado
Department of Computer Science
Boulder, CO 80309
{neufeldm,grunwald}@cs.colorado.edu

CU-CS-971-04 May 2004

University of Colorado at Boulder

Technical Report CU-CS-971-04
Department of Computer Science
Campus Box 430
University of Colorado
Boulder, Colorado 80309

Deafness and Virtual Carrier Sensing with Directional Antennas in 802.11 Networks

Michael Neufeld and Dirk Grunwald
University of Colorado
Department of Computer Science
Boulder, CO 80309
{neufeldm,grunwald}@cs.colorado.edu

May 2004

Abstract

Inexpensive analog phase array antennas are on the verge of becoming widely available. These versatile antennas are capable of very rapidly altering their gain pattern to form complex patterns. However, previous research has shown that problems arise when using the stock 802.11 MAC protocol with directional antennas, and new MAC protocols have been designed to address these issues as well as exploit some of their new capabilites. Unfortunately, most of these protocols sacrifice interoperability with existing 802.11 equipment, making incremental deployment difficult. Even when directionally-aware protocols *are* interoperable with existing equipment, unanticipated problems may arise. In particular, we have found that the problem of *deafness* [1] to be of high importance when directional antennas are used in community networking scenarios.

In this work we present a taxonomy of simple directional enhancements to the 802.11 MAC which maintain interoperability with existing equipment. We also evaluate two of these schemes in various community networking scenarios. In addition, we enhance their resistance to excessive RTS/CTS messages caused by deafness using two different schemes: RTS Validation and RTS/CTS Filtering. Furthermore, we introduce an antenna steering heuristic which trades off some spatial reuse for a decrease in deafness.

1 Introduction

Ad hoc wireless networks are used for a combination of mobile computing and fixed wireless applications. Fixed wireless installations, such as community networks, infrastructure for rural communities or temporary military networks, suffer from limited scalability when commodity network standards are used with omnidirectional antennas [2]. Commercial scalable wireless networks are typically constructed using a combination of directional and omnidirectional antennas arranged to increase spatial reuse of the wireless media. Fixed directional antennas are inexpensive to purchase, but the time and effort required to place and aim them appropriately is relatively high. This makes them less desirable for temporary networks used by first responders or military applications.

An alternative to fixed antennas are “steerable” phase array antennas. These antennas can change orientation or pattern electronically. Inexpensive electronically steerable phased array antennas for wireless data networks are on the verge of being widely available. However, it isn’t immediately clear how best to exploit their capabilities. Previous research using directional antennas has shown that problems arise when using a directional antenna with a MAC layer designed for omnidirectional antennas. Some research has been done on designing new MAC protocols to avoid the pitfalls and take advantage of the features of new directional antennas. Unfortunately, for the most part these new protocols are *not* intended to interoperate with existing 802.11 equipment, making incremental deployment difficult. Furthermore, most of these protocols have been designed and evaluated with mobile *ad hoc* networks in mind. These networks typically have nodes in steady motion and communicating with random peers in the network. These traffic and motion patterns are not representative of traffic conditions arising in fixed wireless networks. This is important since network topology and traffic patterns can have a large effect on network performance, and directional antennas are fundamentally more sensitive to topological conditions than omnidirectional antennas. An example of this sensitivity is *deafness*[1]. This is a significant problem which can occur when using directional antennas.

1.1 Deafness

In Figure 1(a) node A is attempting to send packets to node B. However, node B has its antenna pointed at node C while transmitting to it and is unable to receive packets from A. This will result in excessive timeout and retry

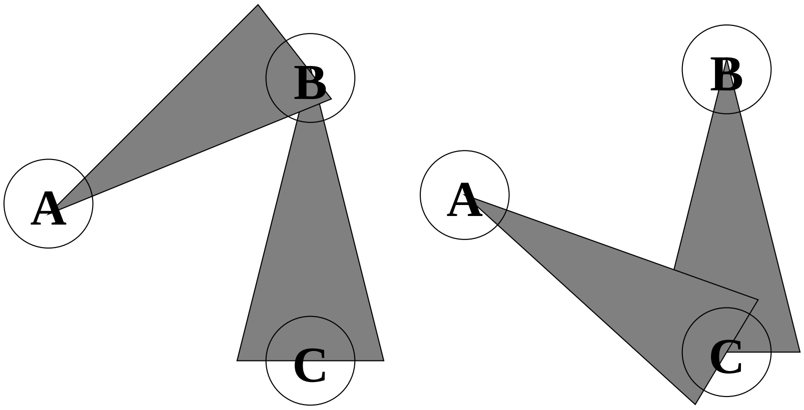

(a) Unheard Packets (b) Colliding Packets

Figure 1: Deafness

at node A. With omnidirectional antennas, node A would have overheard node B using the network and deferred transmission until B became available. Unidirectional transmission has decreased the amount of RF interference generated per transmission, but has also *increased* the number of dropped packets and retransmissions required. Another problem is illustrated in Figure 1(b). In this instance, nodes A and B are attempting to communicate with node C which has its antenna set to an omnidirectional profile. Since A and B can't hear each other's signal they succeed only in colliding with each other. Once again, spatial reuse has been improved at the cost of increased packet retransmission. This is fundamentally similar to the well-known *hidden terminal* problem which also occurs with omnidirectional antennas. The occurence of this problem is highly dependent on the specific network topology and traffic patterns involved. For example, stations keeping up a steady rate of motion and communicating with a diverse set of other stations may not experience this problem to the same extent as static clients connecting to a central internet uplink.

A number of solutions have been proposed to address this issue, *e.g.* adding sideband information channels to inform neighbors when a node is busy. Most of these solutions are not explicitly interoperable with stock 802.11 equipment, and may require additional radio equipment or modifications to the 802.11 physical layer. The basic 802.11 MAC layer does include a mechanism for handling hidden terminals which can also serve to mitigate deafness: the RTS/CTS packet exchange. However, when used in this capacity the RTS/CTS exchange can introduce excess unanswered RTS packets into the network. While these packets are relatively short and may not have a significant effect on *physical* carrier sensing, they can have an extensive impact on *virtual* carrier sensing due to exponential backoff at individual nodes and because they force an update to the Network Allocation Vectors (*NAVs*) of neighboring nodes. This can result in unnecessary deferral of transmission and subsequent underutilization of the network. Deafness may be reduced without completely sacrificing spatial reuse by, for example, transmitting RTS and CTS packets omnidirectionally while still transmitting DATA and ACK packets unidirectionally. While this maintains spatial reuse with respect to the physical carrier sense, it still causes problems with the virtual carrier sense. The number of RTS and CTS messages may be reduced, but sending them omnidirectionally can result in them affecting the NAVs of a larger number of nodes than if they'd been sent unidirectionally. Furthermore, since the actual DATA and ACK transmission and reception are being performed unidirectionally many of the stations receiving the RTS and CTS packets could still safely transmit and receive at the same time.

One possibility for reducing the impact of excessive RTS/CTS messages is to simply *ignore* the ones which aren't followed up by an audible DATA/ACK exchange. Ignoring excess RTS messages has been studied in the context of omnidirectional antennas using a technique called *RTS Validation*[3]. In this work, we propose an additional technique for anticipating the validity of RTS/CTS messages based on past observation called *RTS/CTS Filtering*.

It is also possible to minimize the spatial region over which the RTS/CTS messages need to be sent in order to reduce deafness. Instead of simply choosing an omnidirectional or a narrowly unidirectional gain pattern, a transmission profile can be used which covers *only* the stations necessary to eliminate deafness. While this may revert to a simple omnidirectional transmission in some cases, it can generally provide some spatial reuse enhancements over omnidirectional transmission as well as reduce the deafness and ensuing RTS packets caused by simple unidirectional transmission. In this work we propose a scheme which reduces deafness by trading off some spatial reuse: the *N*etwork-*A*ware *A*ntenna *S*teering *H*euristic, or *NAASH*.

1.2 Contributions of this paper

In this work we present a taxonomy of simple directional enhancements to the 802.11 MAC which maintain interoperability with existing equipment. We also evaluate two of these schemes in various community networking scenarios and enhance their resistance to excessive RTS/CTS messages. We also introduce an antenna steering heuristic which trades off some spatial reuse for a decrease in deafness.

In the next section, we describe antenna characteristics and protocolss proposed improving spectrum efficiency using directional antennas. Following this, section 3 details the protocols we propose. We then evaluate those protocols in §4 and summarize this paper in §5.

2 Prior Work

Before discussing the different algorithms used to control antennas in ad hoc wireless networks, it is useful to understand the different antennas, their capabilities and characteristics. We follow with a brief survey of current research in antennas steering algorithms.

2.1 Types of Directional Antennas

The class of directional antennas encompasses a wide range of devices. While all have the common ability to use directional gain profiles, the specific aiming and steering capabilities vary widely from type to type.

Fixed Beam A fixed beam antenna has a fixed gain profile with a primary lobe pointing in a single direction. Steering the beam is only possible by physically changing the orientation of the antenna, which permits relatively slow changes at best. Quite frequently these antennas are mounted in such a way which allows *no* automatic steering whatsoever. While not particularly flexible, these antennas do provide high gain for their cost and are widely deployed in practice.

Figure 2(a) shows a commercially available parabolic grid unidirectional antenna, and Figure 2(b) a typical gain pattern.

Fixed beam antennas come in many other styles and form factors. For example, flat "patch" antennas may be easily hung on walls. Fixed unidirectional antennas may even be constructed from inexpensive found objects

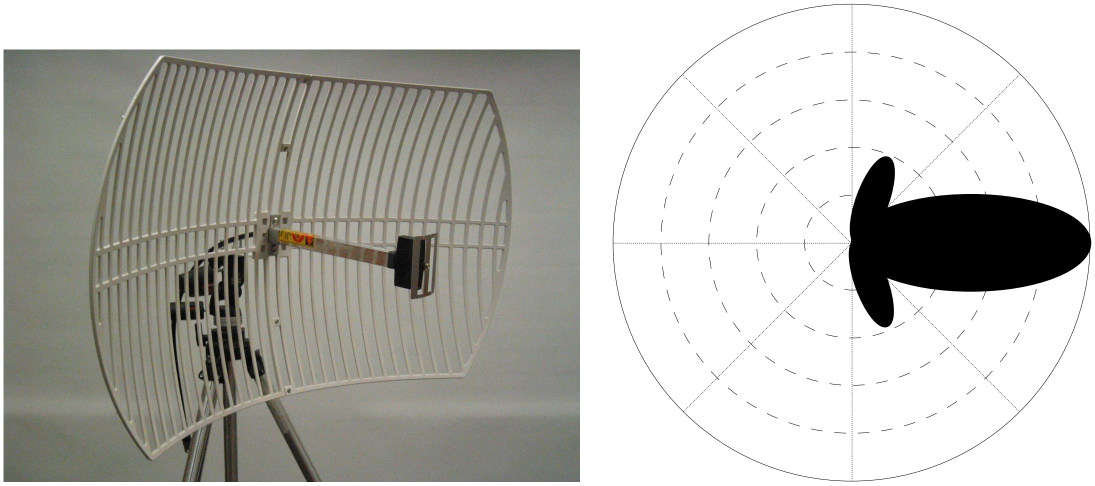

(a) Parabolic Grid Antenna (b) Unidirectional Gain Pattern

Figure 2: Fixed Beam Antenna

such as soup or potato chip cans. The relative inflexibility of this antenna makes deployment and reconfiguration difficult and time-consuming.

Sector A sector antenna consists of multiple fixed beam antennas aimed in different directions, each covering a different sector of space, all of them together giving full 360° coverage. Typcially packets may be sent or received on any *one* sector at a given time. Switching between antennas is done electronically, and is rapid enough to allow the choice of sector to occur on a per packet basis. In practice, these antenna units are often deployed with separate radios for each subantenna, resulting in enhanced spatial reuse when covering a wide area from a single point. Figure 3 shows a gain pattern resulting from four unidirectional antennas. A variation of this antenna type includes a lower gain omnidirectional antenna which may also be used for receving or sending traffic. This allows for a packet to be sent omnidirectionally, albeit at a lower gain than directional traffic.

Analog Phase Array Phase array antennas work by introducing carefully calculated phase shifts into the signal at multiple antenna elements. When combined, the individual signals resulting from these phase shifts interfere constructively and destructively with each other in order to form a particular gain pattern. The phase shift for each element may be changed very rapidly, typically on the order of microseconds. The process of calculating the individual phase shifts for a particular gain pattern may take longer, and depends on the particular physical arrangement of the elements. In practice, a node will likely precalculate phase shifts for particular gain patterns so it may rapidly change between them without excessive overhead.

Figure 4(a) shows an eight element circular arrangement of monopoles connected by phase shifters, and Figures 4(b) and (c) illustrate unidirectional and omnidirectional gain patterns.

Some analog phase array antennas may also provide additional useful information, such as angle of arrival of a signal. However, such features are *not* present in the class of inexpensive analog phase array antennas with which we are primarily concerned.

Digital Phase Array A digital phase array antenna, often referred to as a *smart* antenna uses digital signal processing to accomplish phase shifting. The additional processing power required to do this results in cost and

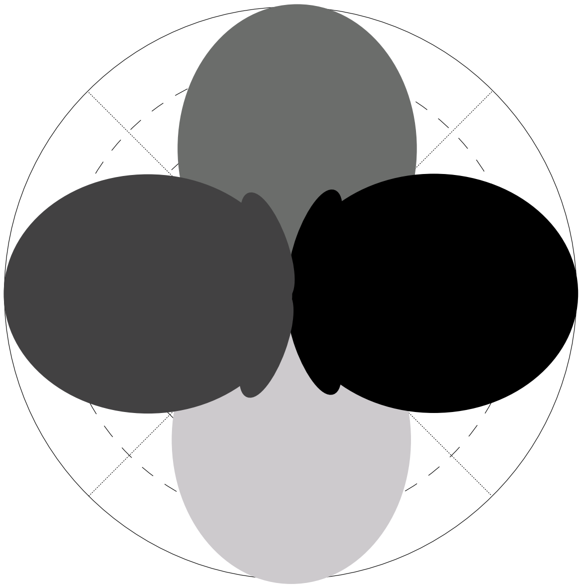

Figure 3: Sector Antenna Gain Pattern

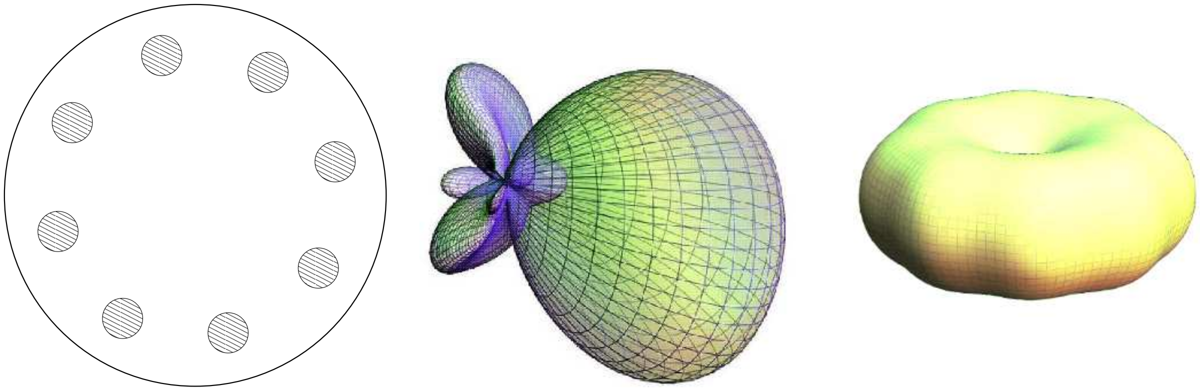

(a) Eight Element Layout (b) Unidirectional Gain Pattern (c) Omnidirectional Gain Pattern

Figure 4: Analog Phase Array - Eight Elements

complexity greater than that of the analog phase array antenna, but also adds significant functionality. Lobes and nulls may be steered very precisely to amplify desired signals and eliminate extraneous ones, and angle of arrival (AoA) information for packets may be obtained as well. With enough processing power, multiple patterns may be realized *simultaneously* using the same set of elements since the phase shift is performed after reception rather than at the elements themselves.

2.2 Antenna Steering MAC Protocols

Much research has been done recently with respect to MAC protocols for directional antennas[4, 5, 6, 1]. Primarily, the approaches taken have been to add directionality to the NAV and to the RTS/CTS exchange, as well as adding side channels (for example *busy tones*) to communicate directional information without consuming bandwidth on the main channel. Some or all neighbors within transmission range are informed of the fact that a particular node is transmitting and isn't available. Additionally there may also be information available about the *direction* in which the node is transmitting, potentially allowing neighboring nodes to achieve better spatial reuse. For the most part, these directional MAC protocols do *not* preserve backward compatibility with stock 802.11 radios, making incremental deployment much more difficult. An exception to this is Directional Virtual Carrier Sensing (DVCS[7]). While it still requires modifying the MAC layer, it is backward compatible with the stock 802.11 MAC. This eases incremental deployment and makes DVCS a more viable deployment option than other directional MAC protocols. With this in mind, we used DVCS as a starting point in order to explore the potential performance benefits of moving to a new MAC layer and more expensive antenna.

Directional Virtual Carrier Sensing DVCS uses an explicitly directional NAV (or *DNAV*) and requires a modified MAC, though it is designed to interoperate with the unmodified 802.11 MAC. It uses the DNAV to maintain information about directional spectrum usage. The RTS/CTS exchange is performed unidirectionally if possible, though it will fall back to using an omnidirectional RTS if node direction is not available or is outdated. DVCS relies on being able to obtain angle of arrival (or *AoA*) information from the antenna, limiting its use to antennas which provide such information, *e.g.* digital phase array antennas. The RTS/CTS exchange is used to determine both direction and signal strength required for sending the data packet. Once both nodes have determined where the other lies, they aim their antennas at each other and perform both transmission and reception unidirectionally. This minimizes both the interference produced by the transmitter and the interference overheard by the receiver.

Circular Directional RTS The Circular Directional RTS scheme[4] is explicitly designed for sector antennas (Section 2.1) where only a single sector may be used for transmitting at a time. It maintains a DNAV to more accurately track spatial RF usage. The RTS packet is sent out sequentially on each sector in a *circular* fashion. Nodes receiving the RTS then note the sector on which the packet arrived, and send out a purely directional CTS on that sector if applicable. The data exchange is then purely directional for both sender and receiver.

ESPAR The ESPAR[5] antenna is an analog phase array. The ESPAR MAC protocol assumes that the antenna is either operating in omnidirectional mode or in a unidirectional mode with a 45° wide lobe aimed in an arbitrary direction. It uses observations of local network traffic, albeit in a different fashion than our heuristic. Instead of using this information to determine the extent of desired local RF coverage, it distributes it throughout the network to allow nodes to make routing decisions which cause the least interference with other traffic streams.

The MAC layer keeps the antenna in omnidirectional mode when idle. When a signal arrives, the receiving antenna does a quick rotational scan to determine the best direction to point to receive the incoming packet. In order to facilitate this scan, nodes transmit a tone signal preceding the data packet which is long enough to cover the scan time. After the scan is complete, the receiving node aims its antenna in the determined best direction. At this point communication occurs unidirectionally.

DMAC/MMAC/ToneDMAC The DMAC [6] protocol uses a combination of a DNAV and a mixture of unidirectional and omnidirectional RTS packets. If the medium is free in all directions, then an omnidirectional RTS is used to inform as many neighbors as possible about a pending transmission. If some directions are blocked by other transmissions, then an unidirectional RTS is used. A multihop RTS extension of DMAC, MMAC [6], has also been proposed.

The ToneDMAC protocol[1] further extends MMAC by utilizing a *busy tone* to to notify neighboring nodes of directional transmissions and deafness. It achieves this by adding a control subchannel to the primary data channel. This channel is used to notify neighbors of directional transmissions (in order to alleviate deafness) without consuming bandwidth on the data channel. It also alters the carrier sense backoff from unidirectional to omnidirectional, eliminating potential deadlock situations due to deafness in multihop situations.

General Directional RTS/CTS There is also some work describing the use of unidirectional and omnidirectional RTS/CTS exchanges without necessarily using a directional NAV [8, 9, 10, 11]. Through both simulation and modelling techniques, this work generally concludes that the best strategy to take is to be as aggressively unidirectional as possible.

2.3 RTS Validation

The RTS Validation technique [3] was originally proposed as a solution to the *blocking* problem occuring when the the omnidirectional 802.11 MAC is used in the presence of hidden terminals. A terminal sending an RTS message suppresses communication between *all* other nodes which hear the RTS, even if that RTS not followed up by a CTS and subseequent data packet exchange. This situation may occur quite frequently in the presence of deafness due to a directional antenna profile. The proposed solution is to *split* the busy wait specified by the RTS into two parts. First, a short wait is performed which allows enough time for the RTS/CTS exchange to occur and the data packet to begin transmission. At that point the physical carrier sense is checked. If the medium is busy, then the remaining time is added to the NAV and deferral continues. However, if the medium is sensed to be idle the remaining wait time is discarded. This allows nodes to ignore RTS messages which are not directly followed up by data transmission.

2.4 Defining Scalability

Prior work typically uses scenarios with point-to-point CBR traffic to evaluate scalability. However, it has been demonstrated that the performance and scalability of a wireless network is dependent on the offered traffic load and topology[12]. The primary focus of our research is the construction of community networks. The point-to-point traffic model is representative of *backbone* traffic providing connectivity over a wide area. Another common traffic pattern in community networks consists of a cluster of nodes communicating with a next hop gateway, *e.g.* a small group of houses sharing a common internet uplink. With that in mind, we use both point-to-point and cluster traffic to evaluate performance.

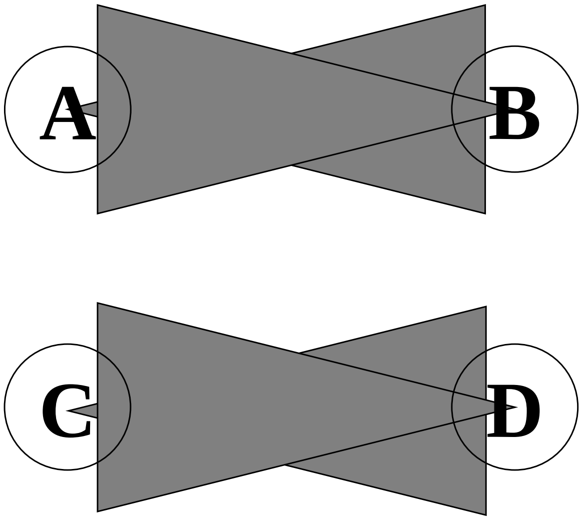

Figure 5: Omnidirectional RTS/CTS Exchange Detrimental to Performance

3 Design

While an aggressively unidirectional protocol is good for enhancing simple spatial reuse, it may also substantially increase the amount of deafness in the network. In order to mitigate this, we propose the use of antenna profiles which are not always unidirectional but still interoperable with the 802.11 MAC. In §3.1 we present a taxonomy of antenna steering schemes which interoperate with the stock 802.11 MAC layer. By explicitly sending RTS and CTS packets to a larger set of stations than the immediate destination, we can reduce deafness. Unfortunately, a broader distribution of RTS and CTS messages can result in unnecessary deferral of transmission and underutilization of the network. For example, Figure 5 shows a case where moving to omnidirectional RTS/CTS packets severely hampers spatial reuse and eliminates the scalability beneifts of directional antennas. Nodes A and B and C and D could safely maintain simultaneous communication if purely unidirectional transmissions are used, but they will needlessly defer to each other if an omnidirectional RTS/CTS exchange is used. In order to both eliminate deafness (as illustrated in Figure 1) as well as maintain spatial reuse, irrelevant RTS and CTS messages must be either be ignored by recipients or never received in the first place. We outlined one technique for ignoring RTS messages, RTS Validation [3], in §2.3. We will also describe a new technique, RTS/CTS Filtering, in §3.2. In order to reduce the number of stations which receive RTS/CTS messages but still reduce deafness, we propose a scheme which tracks the set of stations which could be affected by using a unidirectional antenna orientation. Rather than sending RTS/CTS messages completely omnidirectionally, stations may then opt to cover only that set of stations. This scheme, the **N**eighbor **A**ware **A**ntenna **S**teering **H**euristic (*NAASH*), is described in §3.4.

3.1 Simple Backward-Compatible Directional Antenna Usage

A simple method for ensuring interoperability with the stock 802.11 MAC while taking advantage of directional antennas is to keep the packet format, exchange steps and timings the same, but to transmit and listen for some or all of the packets directionally. Tables 1 and 2 illustrate possible combinations at sending and receiving stations. Note that the last half of the combinations at the receiver, *i.e.* those specifying unidirectional reception of the RTS packet, are only feasible with some preknowledge that a packet is going to be sent.

Short Name	**Send RTS**	**Wait for CTS**	**Transmit DATA**	**Wait for ACK**
Omni	Omnidirectional	Omnidirectional	Omnidirectional	Omnidirectional
Omni/UniRxACK	Omnidirectional	Omnidirectional	Omnidirectional	Unidirectional
Omni/UniTxDATA (OmniRTS)	Omnidirectional	Omnidirectional	Unidirectional	Omnidirectional
Omni RTS/CTS Uni DATA/ACK	Omnidirectional	Omnidirectional	Unidirectional	Unidirectional
Omni/UniRxCTS	Omnidirectional	Unidirectional	Omnidirectional	Omnidirectional
TxOmni/RxUni	Omnidirectional	Unidirectional	Omnidirectional	Unidirectional
Omni 7/10 Split	Omnidirectional	Unidirectional	Unidirectional	Omnidirectional
Uni/OmniTXRTS (Uni w/OmniRTS)	Omnidirectional	Unidirectional	Unidirectional	Unidirectional
Omni/UniTxRTS	Unidirectional	Omnidirectional	Omnidirectional	Omnidirectional
Uni 7/10 Split	Unidirectional	Omnidirectional	Omnidirectional	Unidirectional
TxUni/RxOmni	Unidirectional	Omnidirectional	Unidirectional	Omnidirectional
Uni/OmniRxCTS	Unidirectional	Omnidirectional	Unidirectional	Unidirectional
Uni RTS/CTS Omni DATA/ACK	Unidirectional	Unidirectional	Omnidirectional	Omnidirectional
Uni/OmniTxDATA	Unidirectional	Unidirectional	Omnidirectional	Unidirectional
Uni/OmniRxACK	Unidirectional	Unidirectional	Unidirectional	Omnidirectional
Uni	Unidirectional	Unidirectional	Unidirectional	Unidirectional

Table 1: Potential Directional Antenna Settings for MAC States at Sender

Short Name	**Wait for RTS**	**Send CTS**	**Wait for DATA**	**Send ACK**
Omni	Omnidirectional	Omnidirectional	Omnidirectional	Omnidirectional
Omni/UniTxACK	Omnidirectional	Omnidirectional	Omnidirectional	Unidirectional
TxOmni/RxUni	Omnidirectional	Omnidirectional	Unidirectional	Omnidirectional
Uni/OmniTxCTS	Omnidirectional	Omnidirectional	Unidirectional	Unidirectional
Omni/UniTxCTS	Omnidirectional	Unidirectional	Omnidirectional	Omnidirectional
TxUni/RxOmni	Omnidirectional	Unidirectional	Omnidirectional	Unidirectional
Complete Omni 7/10 Split	Omnidirectional	Unidirectional	Unidirectional	Omnidirectional
Uni	Omnidirectional	Unidirectional	Unidirectional	Unidirectional
Idle Uni/Omni	Unidirectional	Omnidirectional	Omnidirectional	Omnidirectional
Complete Uni 7/10 Split	Unidirectional	Omnidirectional	Omnidirectional	Unidirectional
Complete TxOmni/RxUni	Unidirectional	Omnidirectional	Unidirectional	Omnidirectional
Idle Uni/Uni/OmniTxCTS	Unidirectional	Omnidirectional	Unidirectional	Unidirectional
Idle Uni/Omni/UniTxCTS	Unidirectional	Unidirectional	Omnidirectional	Omnidirectional
Idle Uni/TxUni/RxOmni	Unidirectional	Unidirectional	Omnidirectional	Unidirectional
Idle Uni/Uni/OmniTxACK	Unidirectional	Unidirectional	Unidirectional	Omnidirectional
Pure Uni	Unidirectional	Unidirectional	Unidirectional	Unidirectional

Table 2: Potential Directional Antenna Settings for MAC States at Receiver

Of these potential combinations, we experimented with the aggressively unidirectional configuration (*Uni*) and the unidirectional configuration with omnidirectional transmission of RTS and CTS packets (*Uni/OmniTXRTS* at the sender and *Uni/OmniTxCTS* at the receiver). This is similar to the DVCS [7] protocol, but *not* identical since we did not utilize a DNAV [6].

3.2 RTS/CTS Filtering

RTS/CTS Filtering involves selectively ignoring NAV updates specified by RTS and CTS packets. It does this by tracking communicating pairs of nodes and simply observing which RTS and CTS transmissions have been followed up by corresponding DATA and ACK packets. In order to determine if a particular RTS or CTS between a communicating pair of nodes should be used to update the NAV, it uses a simple *saturating counter*. A saturating counter is a straightforward and efficient mechanism commonly used in microprocessor design to predict whether or not a branch will be taken, and possibly perform some speculative execution based on that prediction. It works by observing past behavior, and incrementing and decrementing a counter. When an event occurs, *e.g.* a conditional branch is taken, the counter is incremented. When an event does *not* occur, *e.g.* a conditional branch is not taken, the counter is decremented. The counter is limited in range between 0 and a specific maximum value, attempting to increment or decrement the counter beyond those ranges has no effect. This keeps the counter responsive as well as limits the number of bits it consumes. In our experiments we set the maximum value to four. This number was chosen arbitrarily, and seemed to perform reasonably well in our simulated environment. We are considering tuning this value in future work. To make a prediction, the value of the counter is checked. If it is at its maximum the prediction is for the event to occur, otherwise it is predicted not to occur. For our purposes, we maintain two separate counters, one for RTS packets, one for CTS. When an RTS or CTS packet is received, the corresponding counter is incremented, receiving a corresponding data or ACK packet decrements the counter. If a counter reaches saturation value for a particular communicating pair of stations, RTS or CTS messages between that pair are *not* used to update the NAV, otherwise they are.

3.3 ***H**ost-based **N**etwork **A**llocation **V**ector: **HNAV***

A potential problem with RTS/CTS Filtering is that it can ignore valid RTS/CTS messages if a different gain pattern is used to transmit them than is used for the actual DATA and ACK packets. For example, Figure 1(b) shows a case where omnidirectional RTS and CTS packets followed up by unidirectional data and ACK packets can cause such problems. Node B hears the RTS/CTS excahnge between nodes A and C, but not the followup DATA/ACK exchange sent unidirectionally. If this occurs enough times, B will start to ignore RTS/CTS exchanges between A and C. So long as B is not attempting to transmit to either A or C, this behavior is desirable. However, if node B attempts to send packets to either of those nodes, it will result in deafness and collision. To avoid this situation, we propose that nodes maintain a separate NAV for each potential network destination. If RTS/CTS Filtering is being employed, this is a trivial addition to make since statistics are already being maintained on a per-pair basis. Upon receipt of an RTS packet, the NAV for the sender of that packet is updated as specified by the duration field. When the corresponding CTS packet is received, the NAVs for both the source and destination are marked as busy. For the virtual carrier sense to indicate an idle channel, *both* the NAV and the HNAV for the destination must be clear.

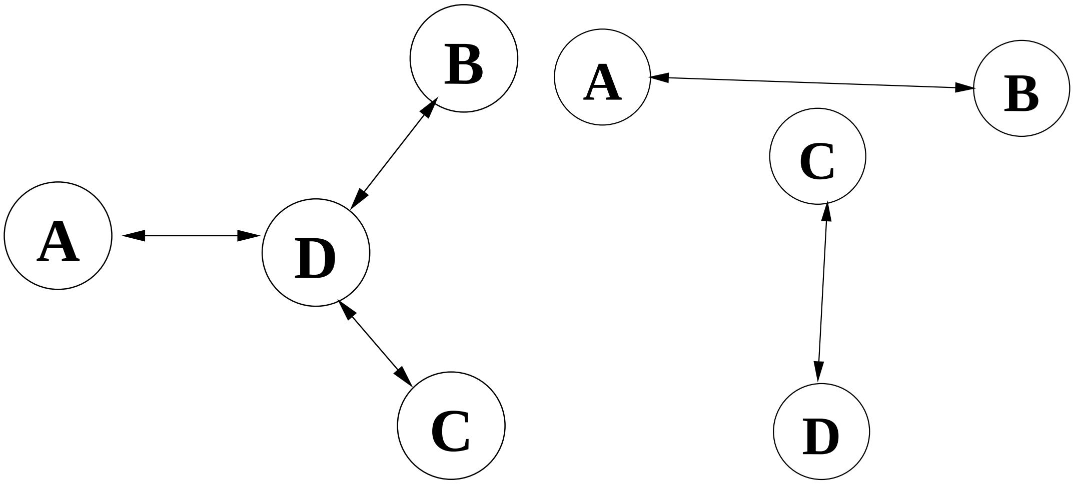

(a) Related Traffic (b) Unrelated Traffic

Figure 6: Related *vs* Unrelated Traffic

3.4 Playing Nice: N*eighbor* A*ware* A*ntenna* S*teering* H*euristic*

Another technique for reducing the effects of deafness and hidden terminals is to reduce the amount of deafness and number of hidden terminals present in the network to begin with. In the trivial limit case, this would amount to giving up spatial reuse and returning to an omnidirectional transmission profile. However, by intelligently observing local traffic patterns it is possible to trade some (but by no means all) spatial reuse for a reduction in deafness and hidden terminals. A key observation to make when considering how to achieve this is that there is a distinction between *related* and *unrelated* traffic. Related traffic contends for the same geographic region, and hence requires cooperation amongst the entities attempting to communicate. Unrelated traffic is between different sets of nodes in different regions. In this case, no cooperation is required and it's best if the traffic simply isn't heard at all. Figure 6(a) shows an example of related traffic. Nodes A,B, and C are all communicating back and forth with node D, which has an uplink to the internet. Figure 6(b) shows two sets of unrelated traffic, as might occur between two separate long haul backbone links. Nodes A and B are communicating with each other, as are nodes C and D. The goal of the Neighbor Aware Antenna Steering Heuristic (NAASH) is to keep track of related traffic and aim the antenna to promote cooperation between nodes when appropriate and to minimize the interference of unrelated traffic. A simple example of this is shown in Figure 7. Nodes A and B are attempting to communicate with C. Figure 7(a) illustrates what happens when they both use narrow unidirectional gain patterns. Since neither can hear the other, the chance of collision at C occuring due to simultaneous transmission is increased. Figure 7(b) shows the results of moving to a somewhat wider transmission beam. Since A and B can hear each other transmitting, they can cooperate to reduce the number of collisions occuring at C. While they have sacrificed some overall spatial reuse, they are communicating more effectively with each other.

NAASH requires two data structures to function: a table of active traffic in the immediate area and angle of arrival estimates for neighbors. The **A***ctive* **T***raffic* **T***able*, or *ATT*, keeps track of who is actively communicating with whom in the local area. For example, if node C overhears node A sending packets to node B, then the ATT at node C would contain an entry for node B indicating that node A had been overheard sending traffic to it. Our implementation also includes a field indicating *when* the last packet was received so that only recent traffic is taken into account. The ATT is kept up to date by promiscuously listening for traffic on the wireless network. When a packet arrives, its source and destination MAC addresses are extracted. The arrival time of this packet

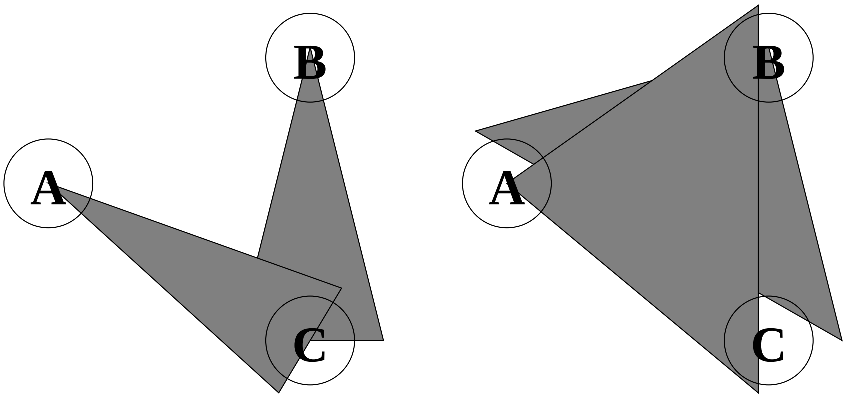

(a) Hidden Terminals, Colliding Packets (b) Hidden Terminals Eliminated

Figure 7: Hidden Terminals Eliminated

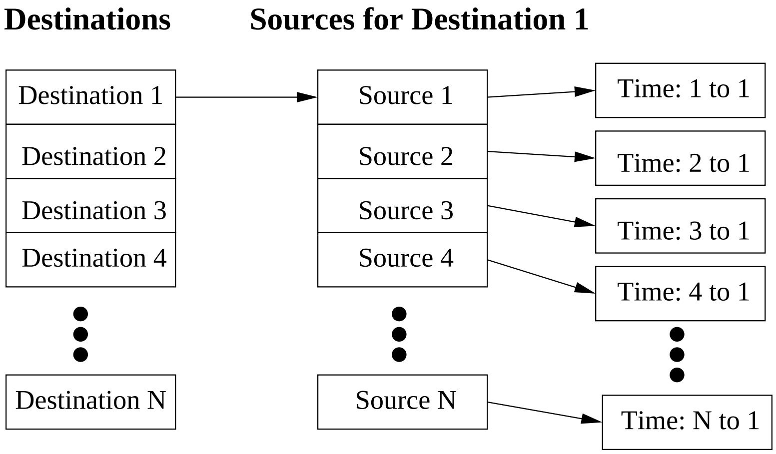

Figure 8: Active Traffic Table

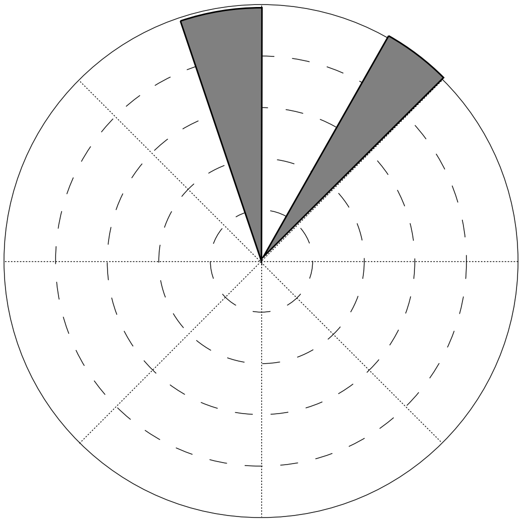

(a) Graphical

0					5					10					15					20					25				
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0

(b) Neighbor Location Table

Figure 9: Two Possible Locations for a Neighbor - $45°$ and $90°$

is then stored, indexed primarily by the destination and secondarily by the source in a two layer hash, illustrated in Figure 8. This structure allows NAASH to determine the set of nodes a particular transmission should cover based on the destination address of a packet by looking up all of the sources which have been observed recently transmitting to a particular destination. Once the set of desired nodes has been determined, NAASH must then configure the antenna to cover them. Doing this requires an estimate of the angles at which those neighbors lie. If the best transmission direction for each neighbor were known exactly, this table could simply be a hash mapping neighbor addresses to angles. However, dealing with potentially inexact location information requires a more complicated structure. In our implementation we use a bitmap containing 360 entries, allowing for multiple, potentially discontinuous ranges of locations to be specified to the nearest degree. This structure, the **N**eighbor **L**ocation **T**able, or *NLT*, is illustrated in Figure 9.

When a unicast packet is queued for sending at a node, NAASH uses the ATT to find the active neighbors for that destination and the node itself. [1] This is the list of nodes which will be covered by the packet transmission. NAASH then takes this list and consults the NLT to find all of the locations which must be covered. These locations are assembled into a single bitmap by performing a logical OR on the location bitmaps for all of the covered nodes. This bitmap is then used to calculate the narrowest transmission cone which will cover all of the node locations. The longest gap in potential locations is found, *i.e.* the longest run of zeros in the bitmap. The edges of this gap are taken as the boundaries of the transmission cone. In the context of Tables 1 and 2 this adds a third possible configuration for the antenna at each entry.

4 Evaluation

We used the Click Modular Router[13] for our implementation and the *nsclick*[14] simulation environment to evaluate our scheme inside of *ns-2*. We added support for directional antennas and MAC layer directional enhancements to *ns-2*.

---

[1] Broadcast packets are treated as a special case and sent omnidirectionally.

4.1 Simulated Scenarios

Our node topology was based on house positions digitized from a TerraServer[15] map of a section of Boulder, CO. The radio parameters used were intended to match a particular variety of Orinoco card operating at 2 Mb/s[16]. The transmission range covers the entire simulation area. Our directional antennas could form a directional 45° transmission cone in unidirectional mode, approximately corresponding to the capabilities of the phase array antenna we plan on using in our future deployment. We assumed that the antennas could provide exact angle of arrival information for arriving packets, though this will not be true for our actual hardware. We studied two major kinds of traffic: groups of nodes communicating with a central server, representing nodes sending traffic to and from an uplink, and randomly chosen pairs of nodes communicating with each other, representing long distance "backbone" connections. Examples of the cluster and backbone scenarios we used are shown in Figures 10 and 11.

Our cluster scenarios were further broken down into two subcategories: multiple relatively lightly loaded clusters, and single variably loaded clusters. The first set of scenarios is intended to gauge wide area scalability between multiple non-cooperating clusters, the second for scalability within a single cluster. For the multiple cluster and backbone scenarios, each stream was CBR traffic, 10 packets per second, 512 bytes per packet. To vary the overall load on the network, we increased the number of clusters. For the backbone connections we varied the load by varying the number of pairs of communicating nodes. For the single cluster scenarios, each stream was CBR traffic 1024 bytes per packet, and the packet rate was varied from two to twelve packets per second. All confidence intervals shown are 95% unless otherwise stated. The single cluster scenarios do not serve to evaluate ability to exploit spatial reuse. The bottleneck in these cases is the single uplink node, and a narrower transmission profile will not increase the oveall bandwidth of this link. However, these scenarios can highlight deafness and hidden terminal problems introduced by directional antennas. Packet routing was performed using ARP broadcast and reply, which is equivalent to using a reactive *ad hoc* routing protocol in a single hop environment.

4.2 Performance Metrics

We looked at two primary metrics when comparing simulation results: simple goodput and latency, and the ratio of RTS messages sent to the number of CBR packets offered for delivery. The use of goodput and latency is straightforward. The ratio of RTS messages to CBR packets offered is useful for gauging spatial reuse *vs* deafness and hidden terminals. A high delivery rate combined with an RTS/CBR ratio of around 1.0 implies that the network is experiencing little to no deafness since RTS messages aren't being retransmitted often and packets are being delivered. As the ratio drifts higher, more RTS messages are being sent per offered packet, implying deafness and congestion. If the ratio is *below* 1.0 then stations are not attempting to send all of their packets and could be spending too much time deferring. This indicates that the antenna steering heuristic is too conservative and is not exploiting as much spatial reuse as it could, or simply that there isn't any spatial reuse to be had.

4.3 Baseline Unidirectional Performance

In this section, we will review the base performance for the purely unidirectional steering protocol. Figures 12 and 13 show performance results for the multiple cluster scenarios. In this case, both RTS Validation and RTS/CTS Filtering show some advantages over the unmodified NAV at higher load, both for simple performance and the quantity of RTS messages sent. An important feature to note is that the confidence intervals get very

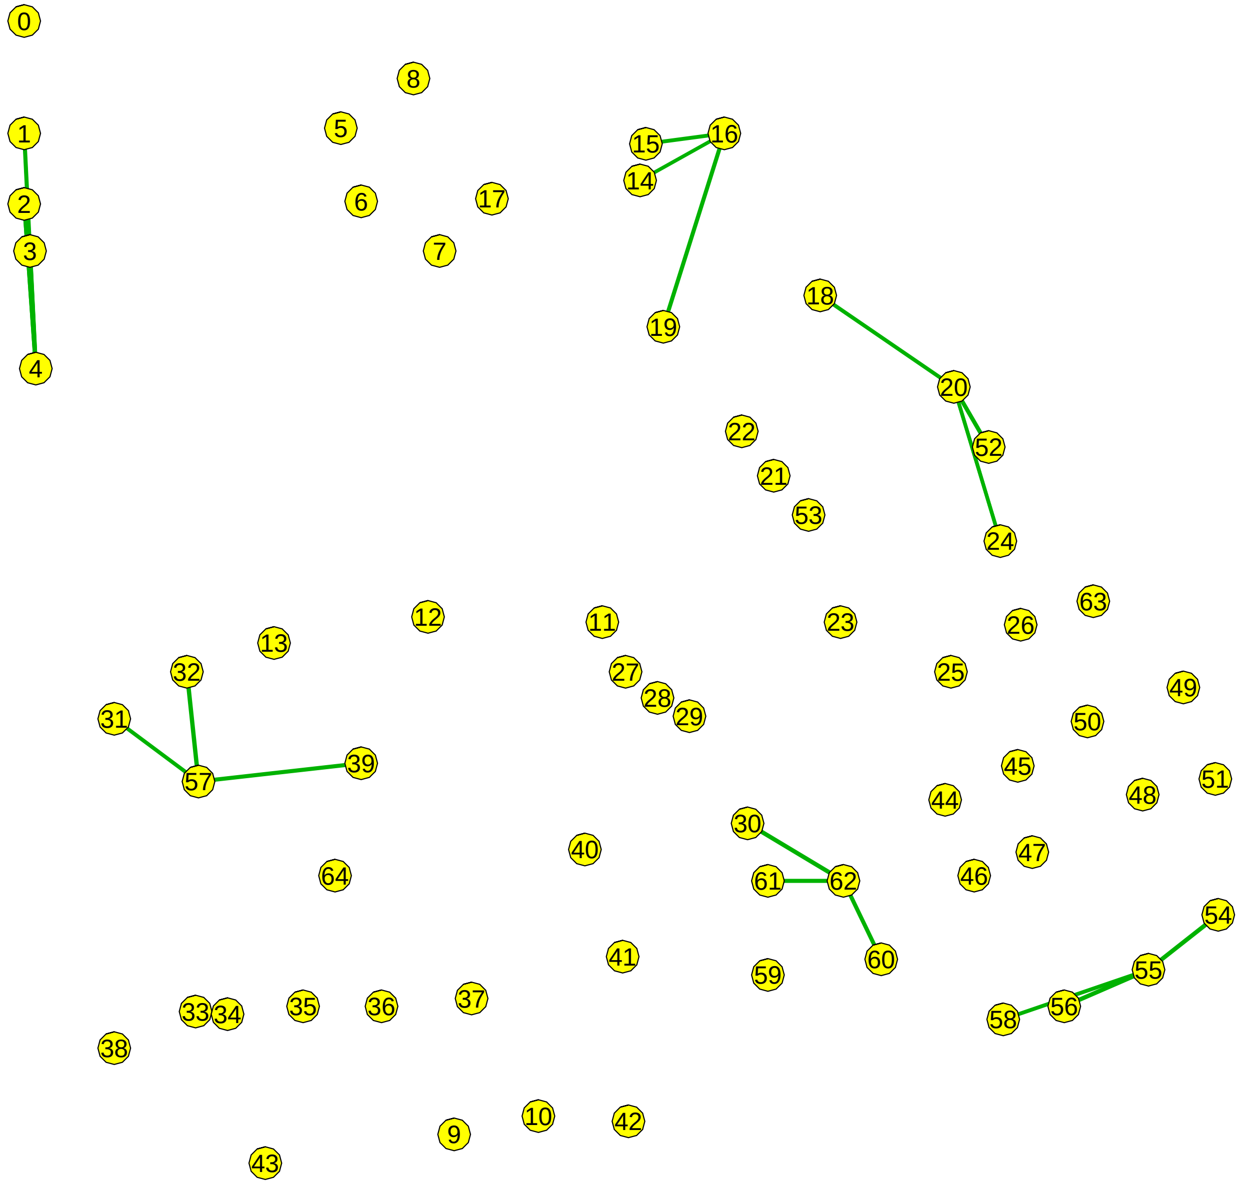

Figure 10: Network Traffic Scenarios: Clusters

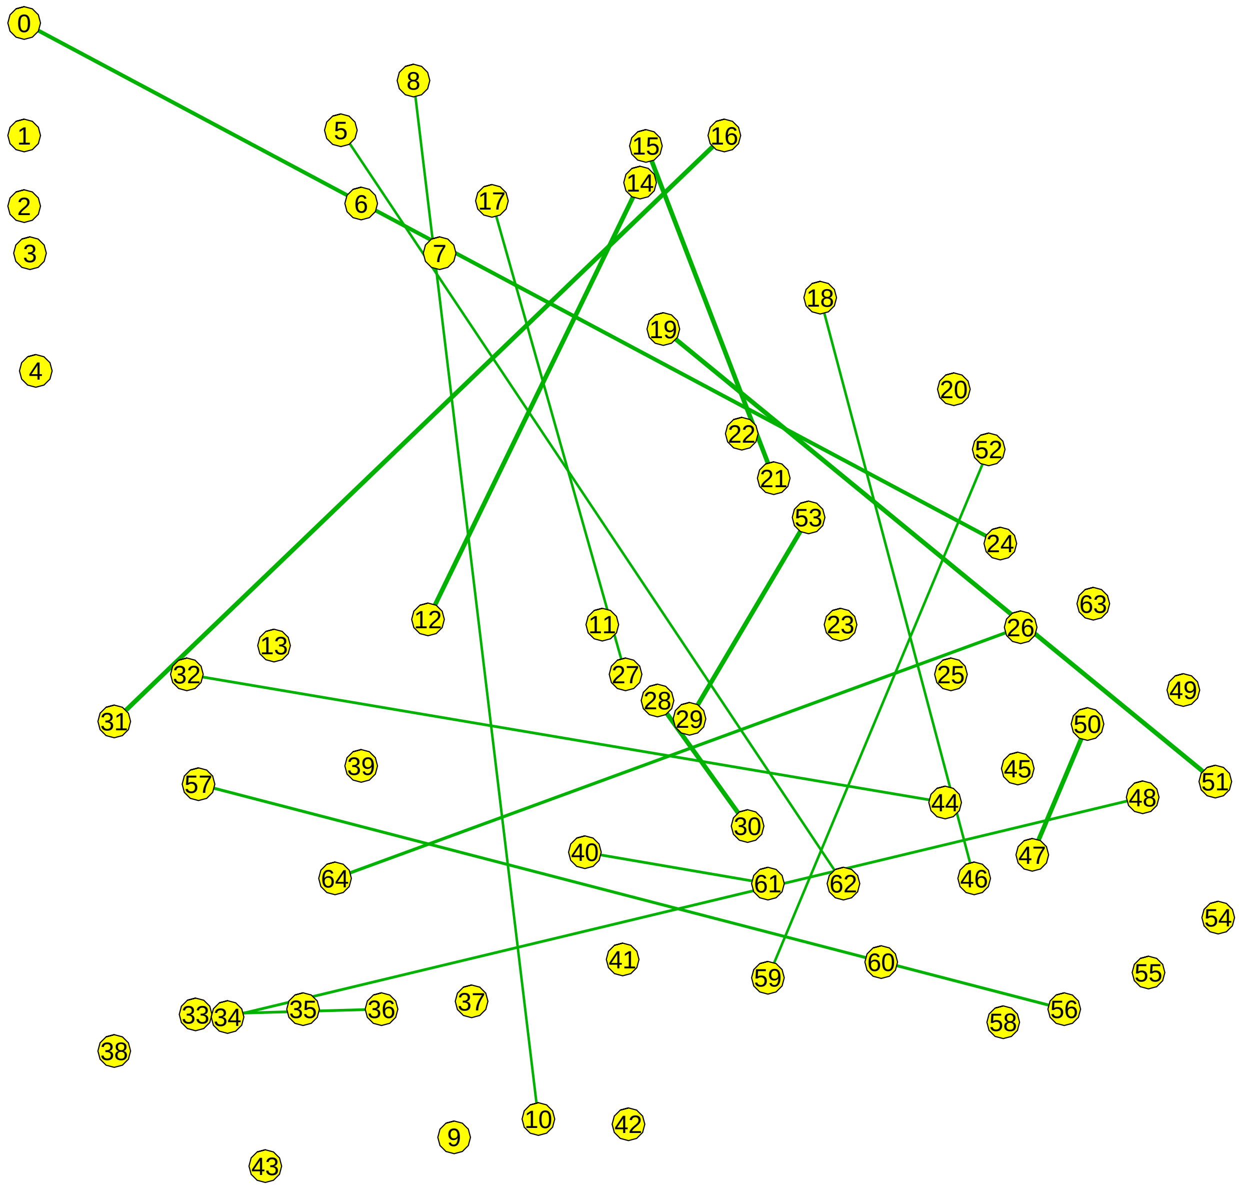

Figure 11: Network Traffic Scenarios: Backbones

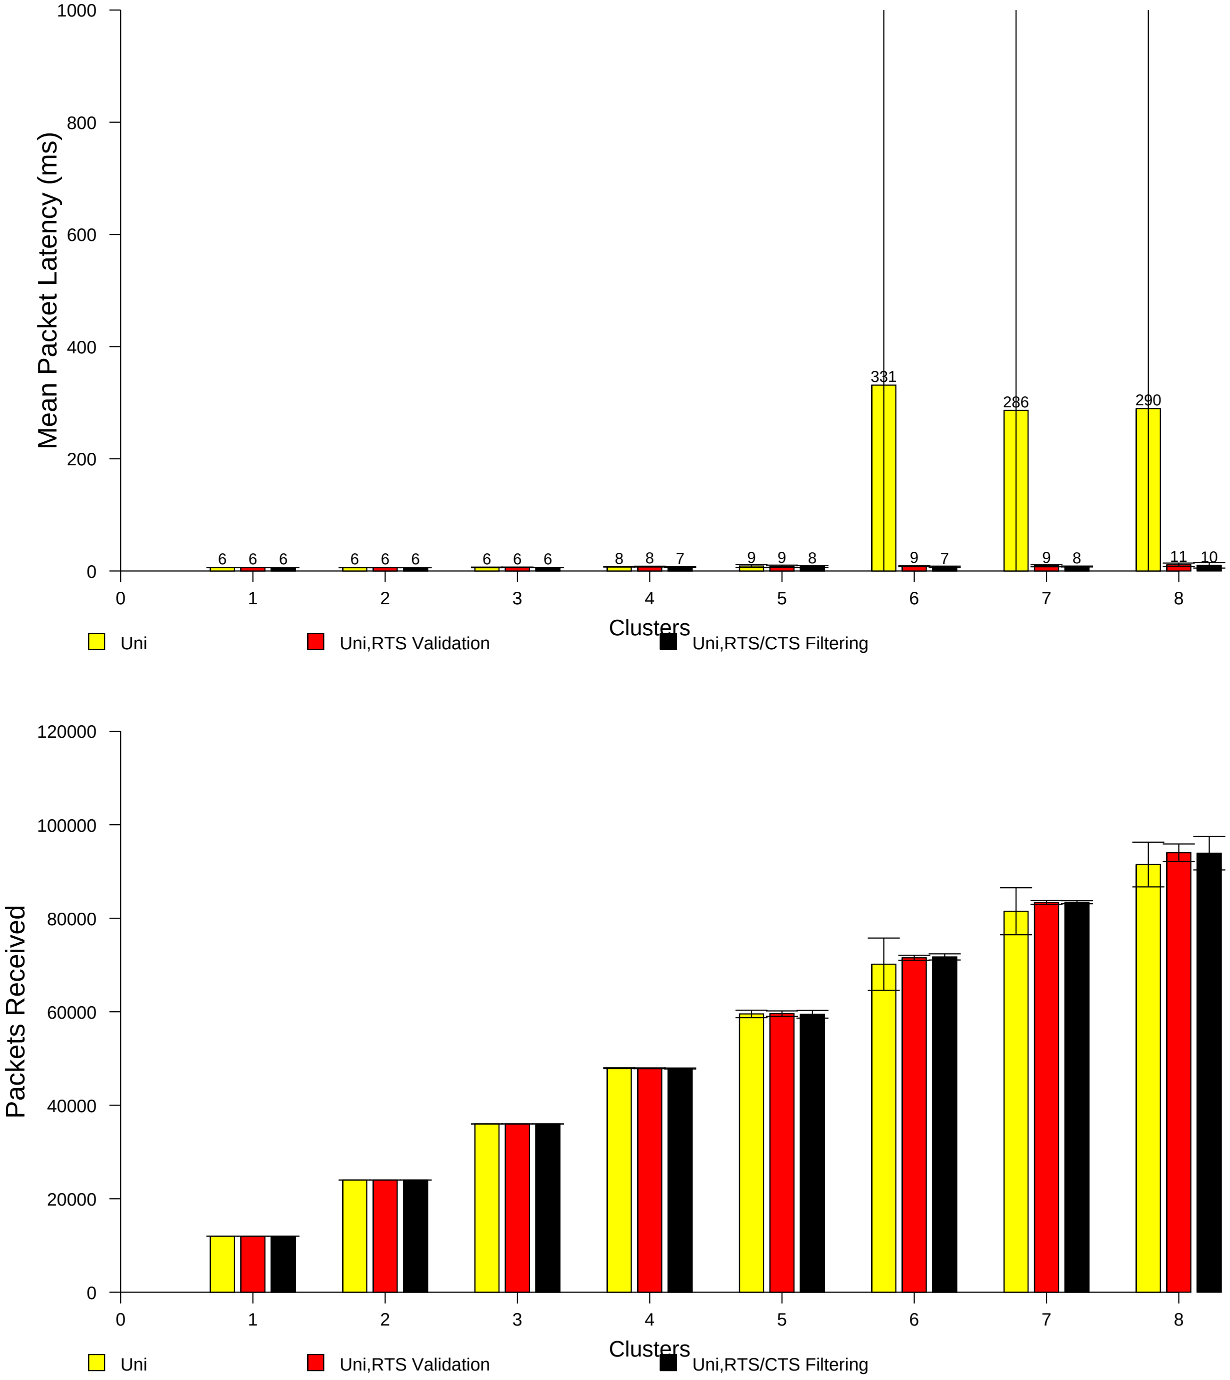

Figure 12: Unidirectional,Multiple Cluster Traffic,Goodput and Latency

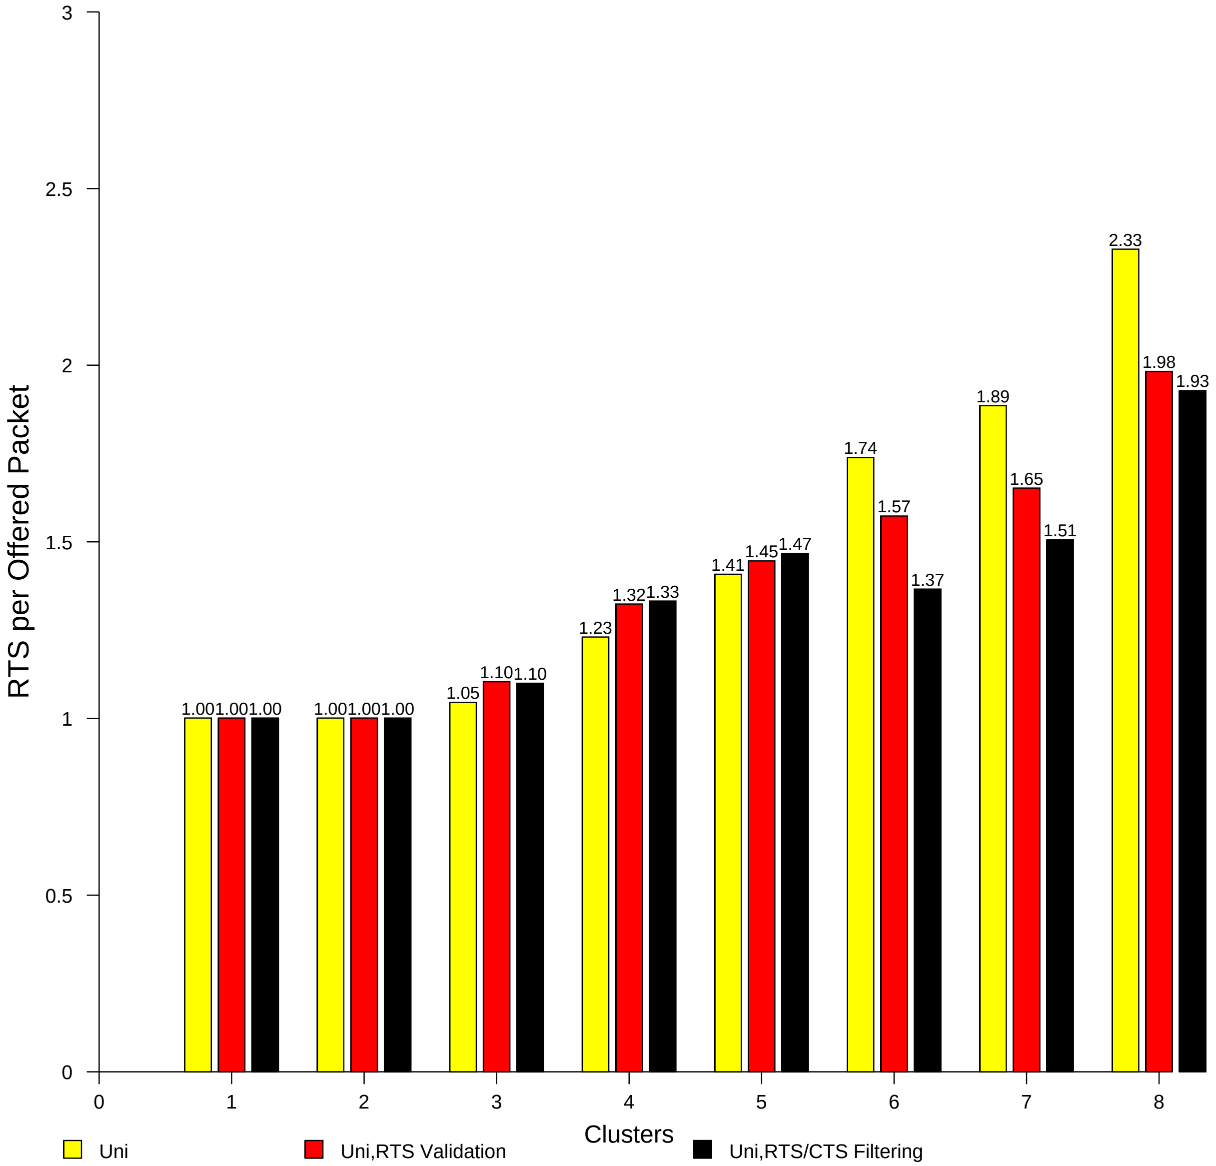

Figure 13: Unidirectional,Multiple Cluster Traffic,RTS per CBR Packet Offered

wide as performance degrades. This is due to unfairness in some scenarios, resulting in a very wide range of performance depending on the specific network configuration. While the mean latency may be high overall, some streams are still achieving very acceptable performance.

Figures 14 and 15 show results for the backbone scenarios. In thse scenarios neither RTS Validation nor RTS/CTS Filtering has an appreciable impact on performance. Figure 15 shows an RTS to CBR ratio consistent with the multiple cluster scenarios in Figure 13.

Finally, Figures 16 and 17 show results for the single cluster configurations. For these, RTS Validation results in a definite performance improvement. RTS/CTS Filtering has a somewhat negative effect, though performance is far from good in either case. Figure 17 shows overall RTS per offered data packet ratios which are much larger than those for the backbone and multiple cluster scenarios. This indicates that the single cluster case may benefit the most from reduced deafness.

4.4 Omnidirectional RTS/CTS

As is shown in Figures 18 and 19 there is some performance benefit to moving to an omnidirectional RTS/CTS exchange for the single cluster scenarios. However, since spatial reuse is *not* a factor in this case it is important to examine results for the cluster and backbone scenarios to see how much spatial reuse has been given up. As shown in Figure 20, performance in the multiple cluster scenario suffers a serious reduction in scalability. RTS/CTS Filtering mitigates it to some extent, but RTS Validation has little to no effect. The reasons for this are indicated by Figure 21. The simple omnidirectional RTS/CTS exchange results in much more conservative usage of the medium, with its RTS per packet ratio dropping well below 1.0. RTS Validation doesn't have a significant effect on this, meaning that it is unable to take advantage of its short deferral period.

The backbone scenarios show very different results. As Figure 22 illustrates, the omnidirectional RTS/CTS exchange results in a reduction in scalability. However, RTS/CTS Filtering has a substantial *negative* effect on top of this. Figure 23 shows that it is too aggressive at ignoring RTS/CTS NAV updates, resulting in a flood of RTS/CTS retries. RTS/CTS Filtering is vulnerable to cases where the overall activity and noise level is high since it may make a decision to ignore RTS/CTS exchanges from a particular pair of nodes due to collisions with traffic from *other unrelated* traffic in the system. An instance of this is illustrated in Figure 24. Stations A and B and C and D are able to safely communicate simultaneously and independently due to their unidirectional antenna configurations. Station E, however, can hear transmissions from both streams. This greatly increases the chances that collisions will occur at E. Since RTS/CTS Filtering treats a failure to overhear a full RTS/CTS/DATA/ACK exchange as an indicator that it should ignore RTS/CTS messages between the pair of stations in question, this may result in stations too aggressively ignoring RTS/CTS messages. In contrast to this, RTS Validation relies generically on the *physical* carrier sense to determine if it should use the entire waiting period specified in the RTS packet. As the network gets busier, the physical carrier sense is busy more often. This results in RTS Validation behaving more conservatively. This is quite evident from both Figures 21 and 23, which show RTS/CTS Filtering steadily becoming more aggressive as the load increases and RTS Validation more conservative.

4.5 NAASH RTS/CTS

In order to further improve the scalability of the omnidirectional RTS/CTS exchange while still decreasing deafness, we also performed experiments where the NAASH antenna profile was used for the RTS/CTS exchange. As is shown in Figures 25 and 26, performance is comparable to using omnidirectional RTS/CTS packets.

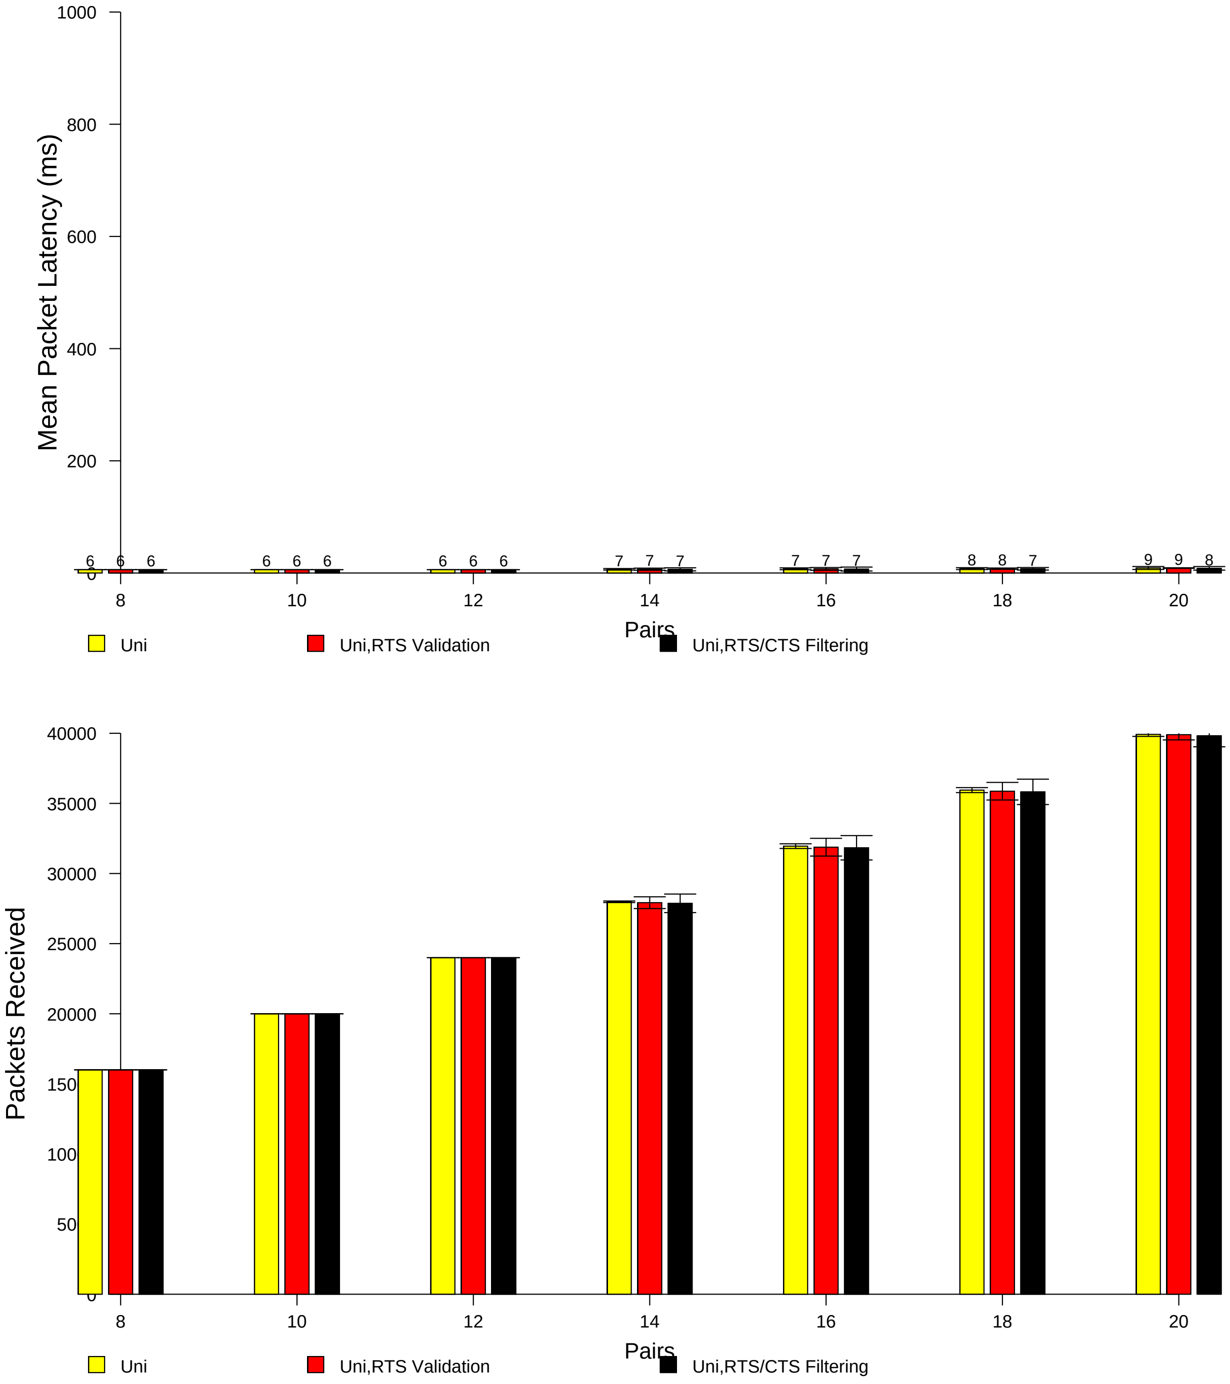

Figure 14: Unidirectional,Backbone Traffic,Goodput and Latency

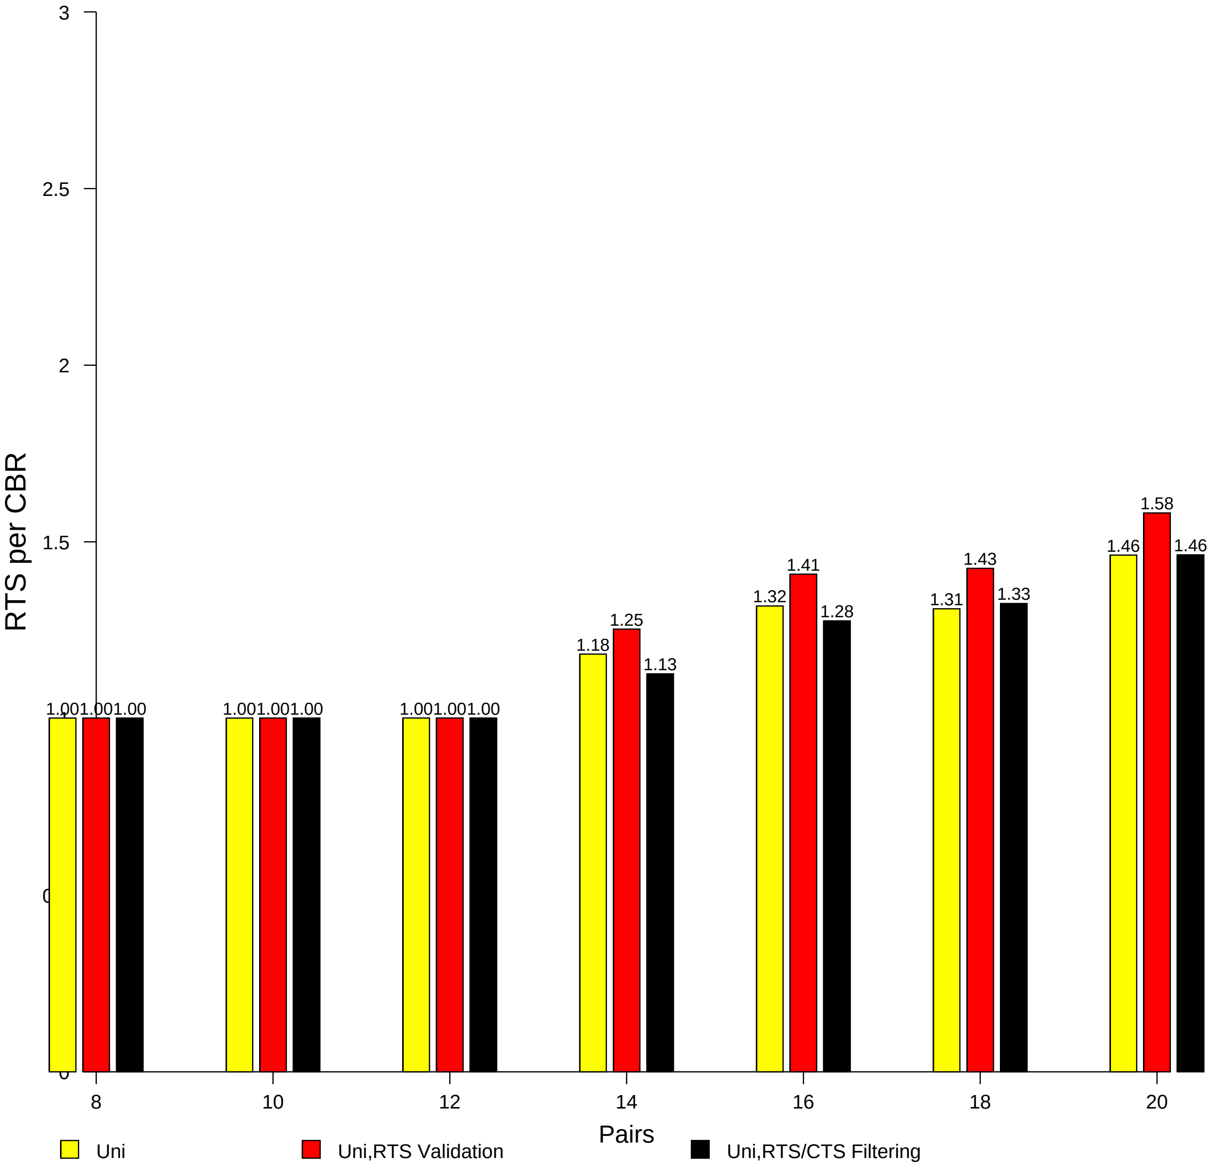

Figure 15: Unidirectional,Backbone Traffic,RTS per CBR Packet Offered

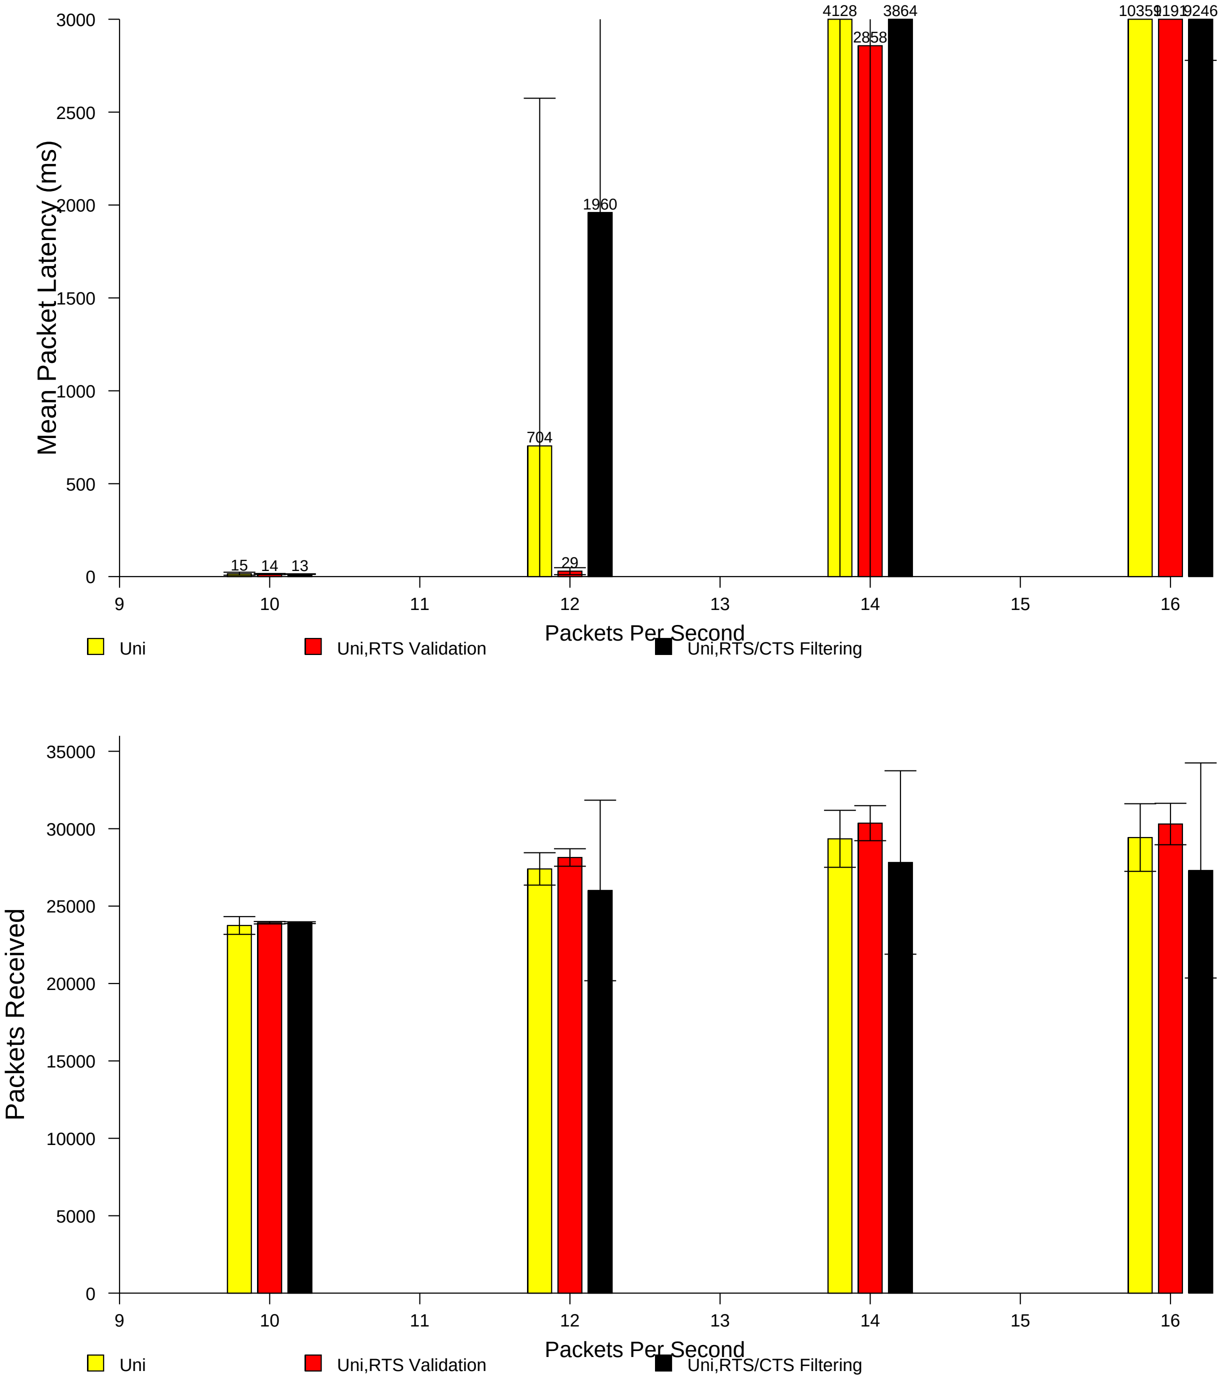

Figure 16: Unidirectional,Single Cluster Traffic,Goodput and Latency

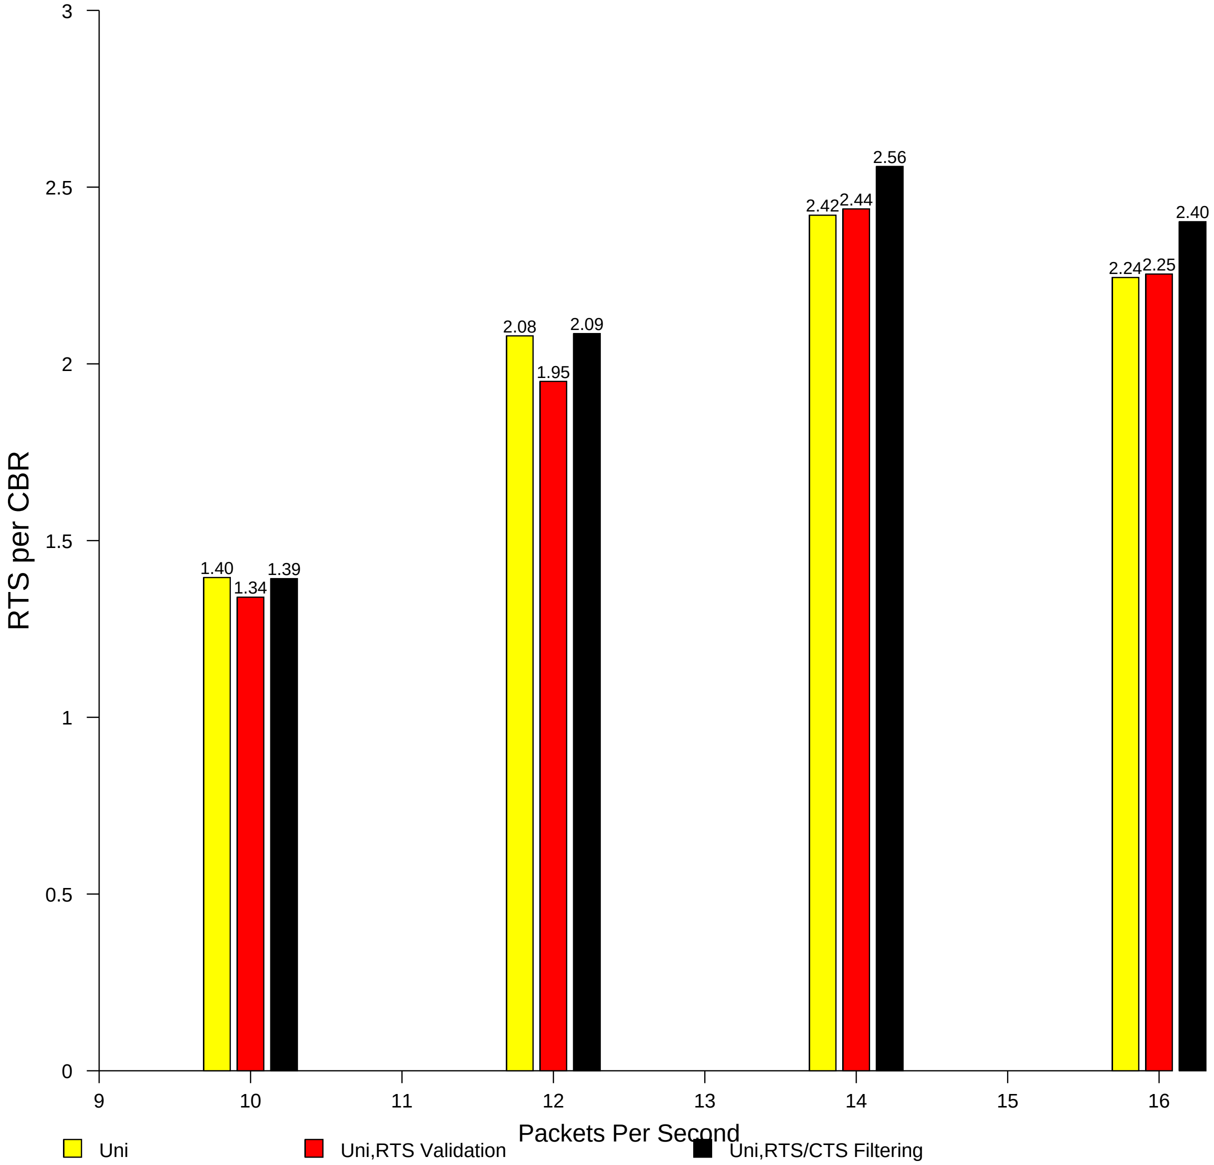

Figure 17: Unidirectional,Single Cluster Traffic,RTS per CBR Packet Offered

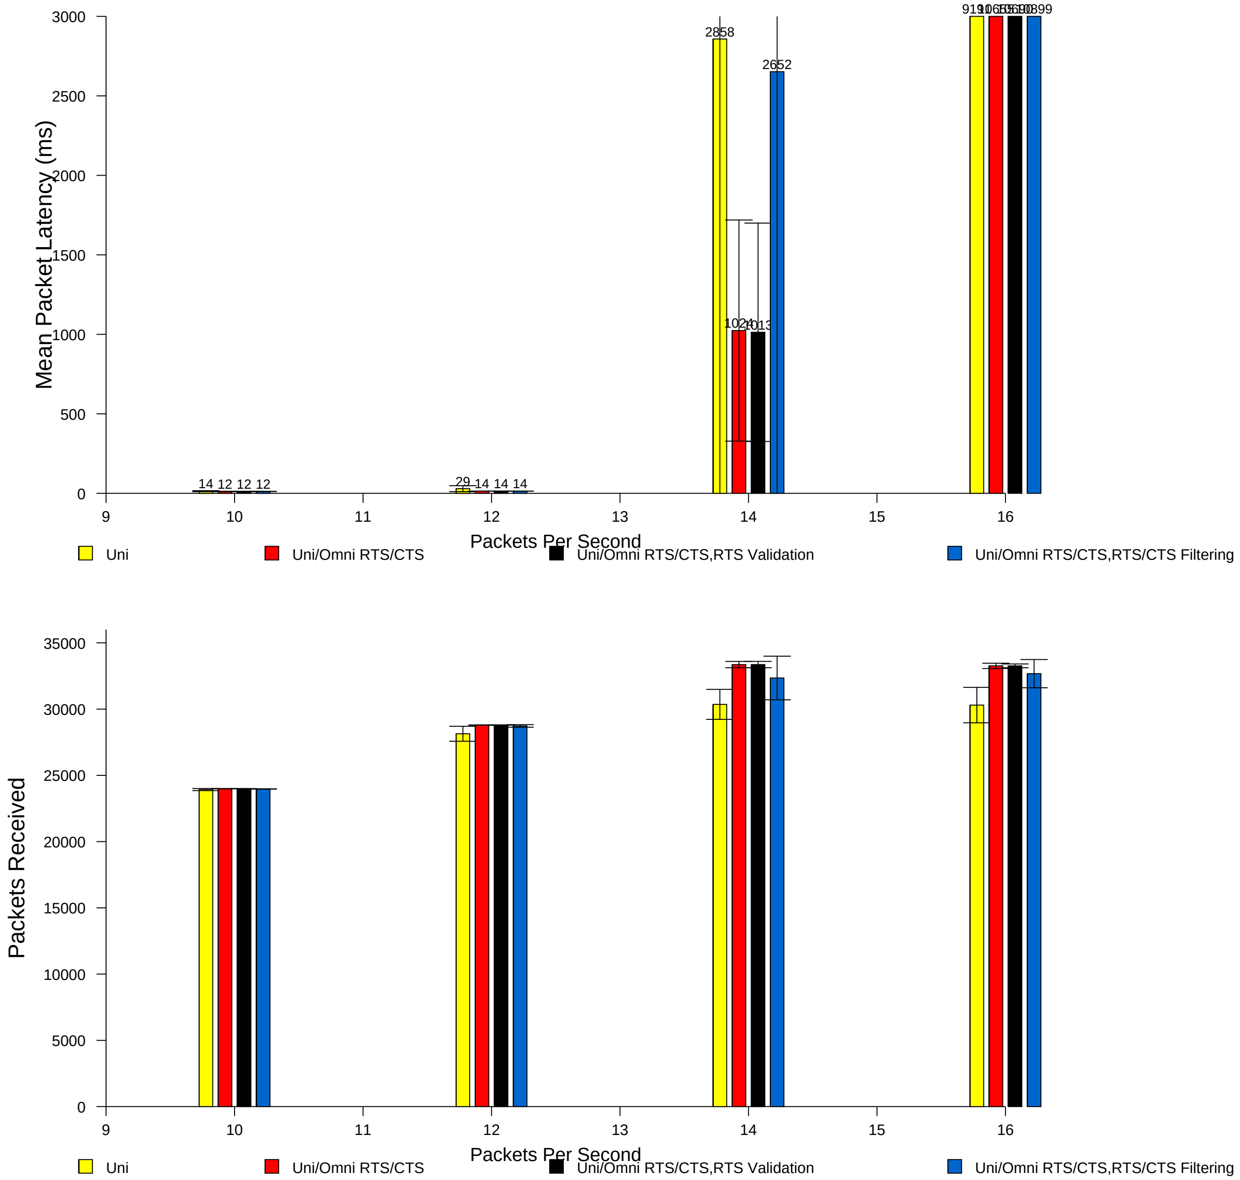

Figure 18: Omni RTS/CTS,Single Cluster Traffic,Goodput and Latency

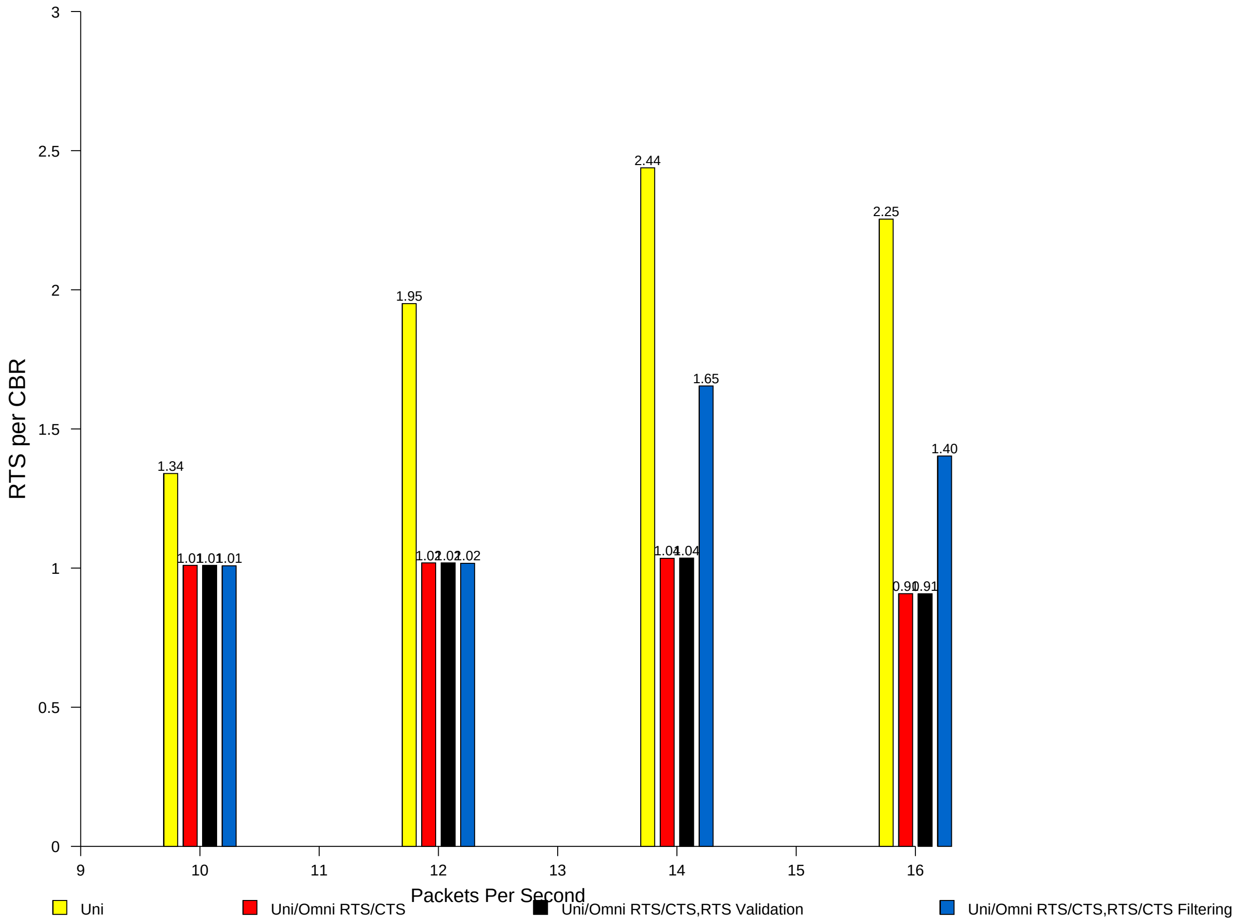

Figure 19: Omni RTS/CTS,Single Cluster Traffic,RTS per CBR Packet Offered

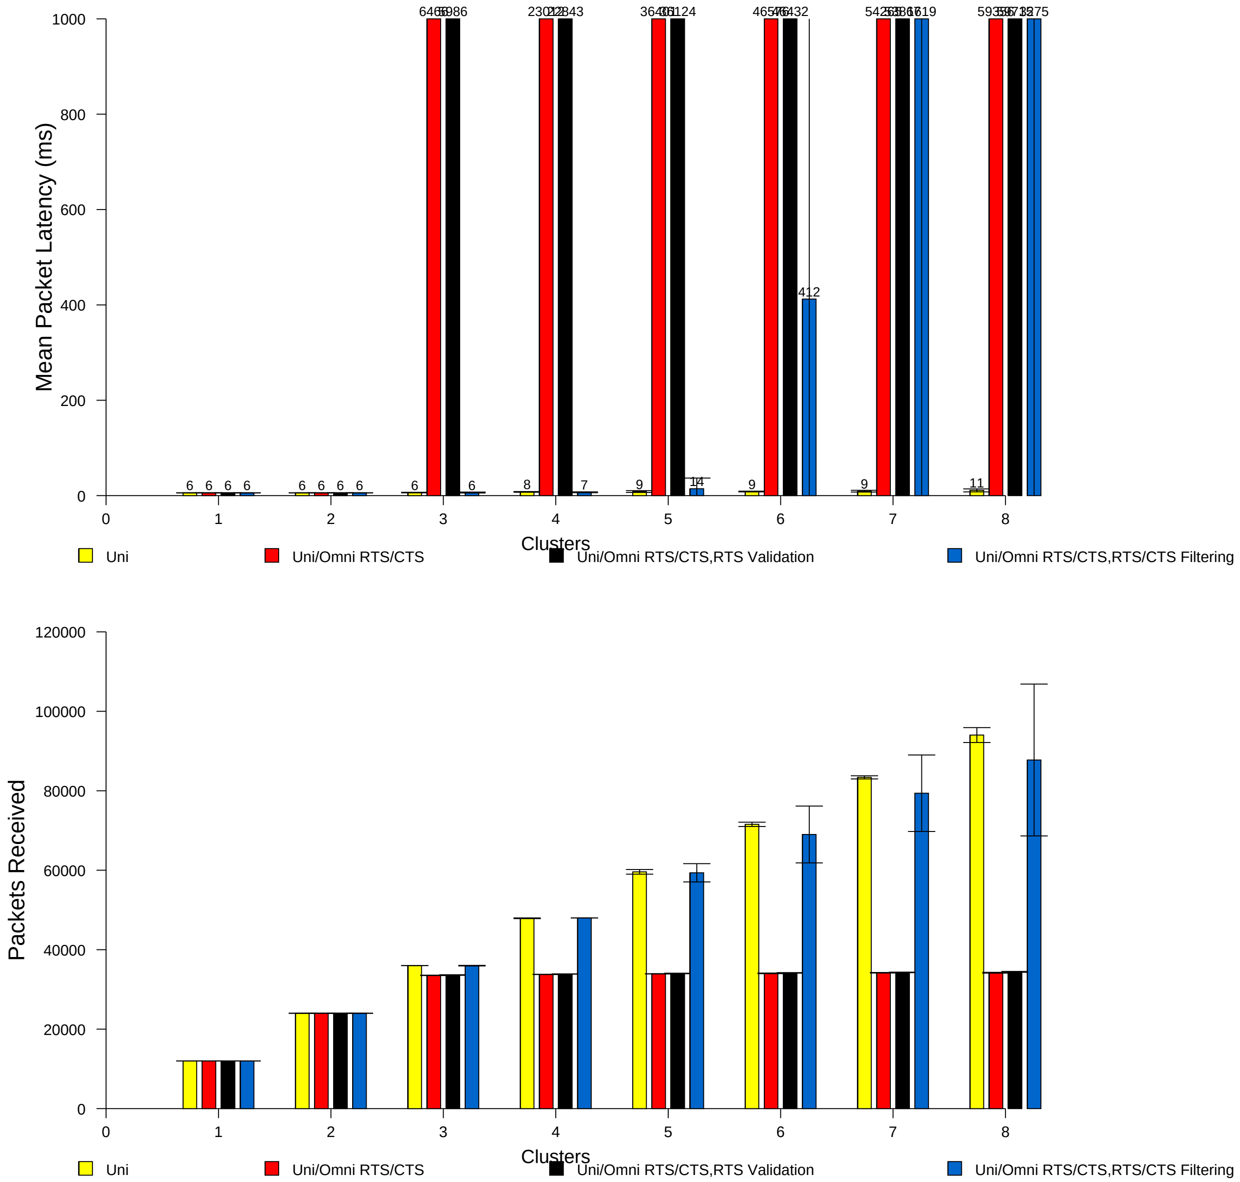

Figure 20: Omni RTS/CTS,Multiple Cluster Traffic,Goodput and Latency

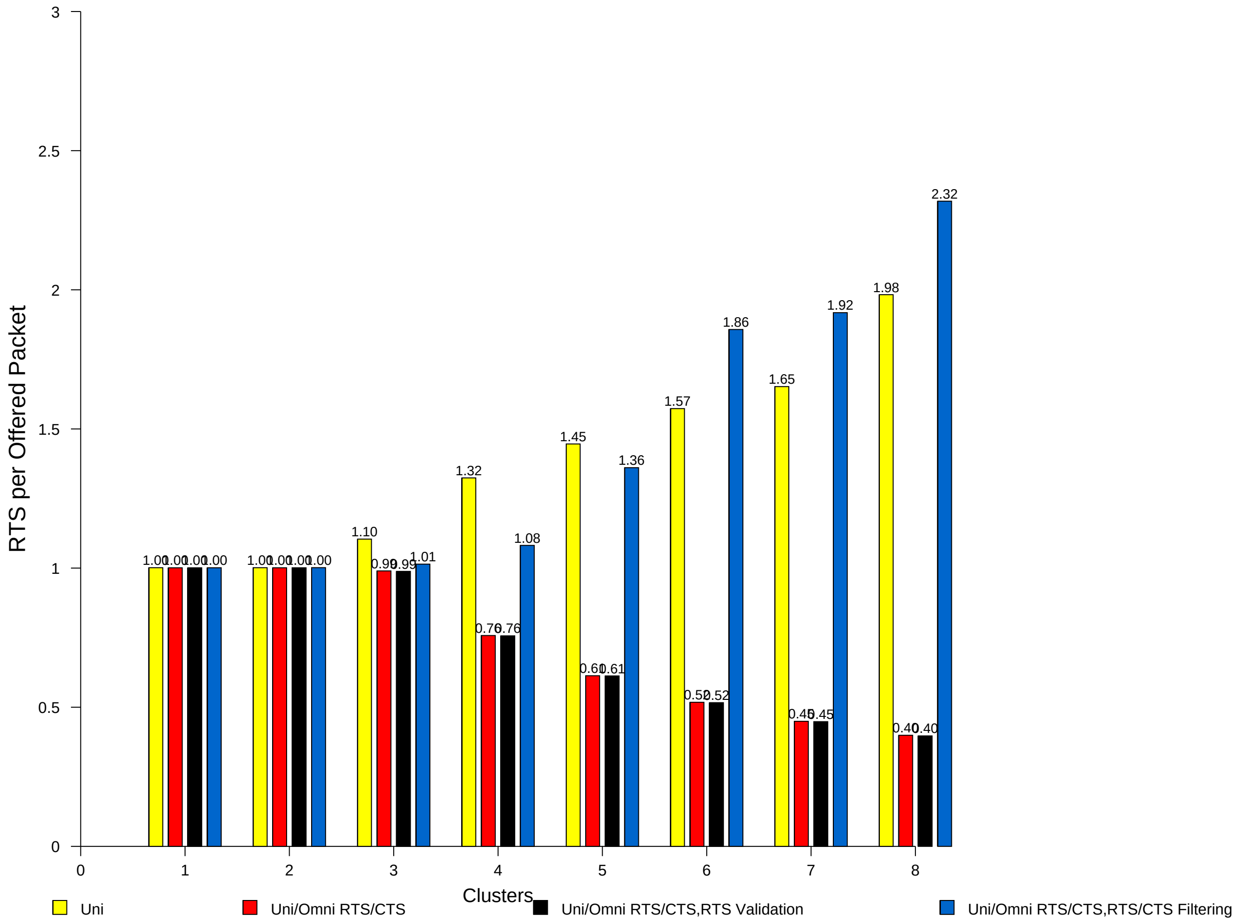

Figure 21: Omni RTS/CTS,Multiple Cluster Traffic,RTS per CBR Packet Offered

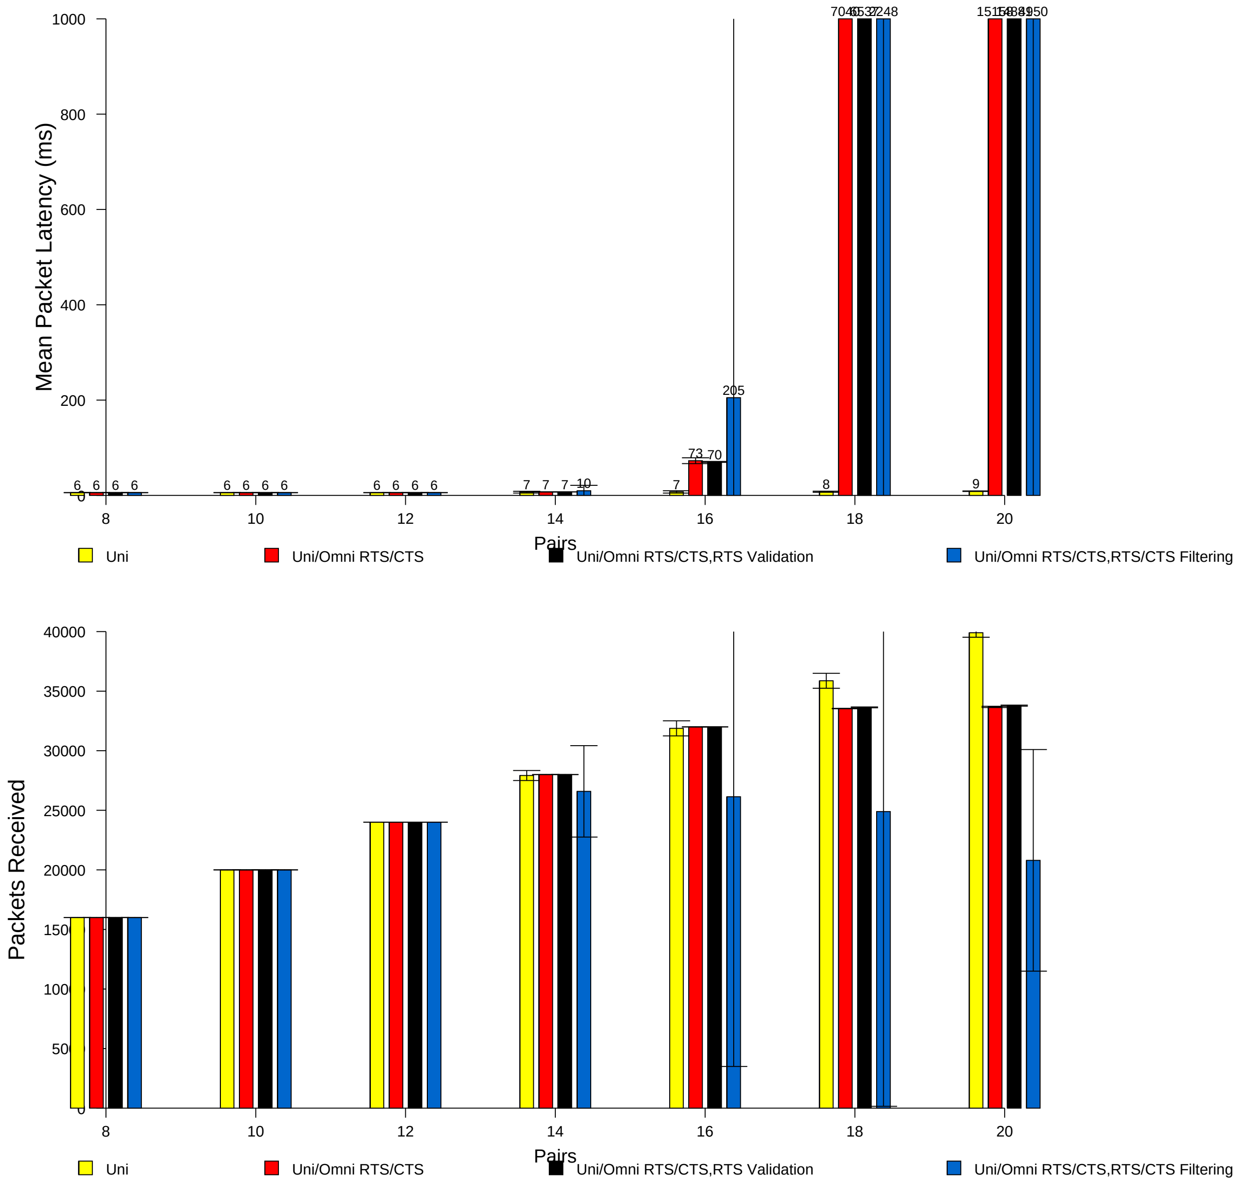

Figure 22: Omni RTS/CTS,Backbone Traffic,Goodput and Latency

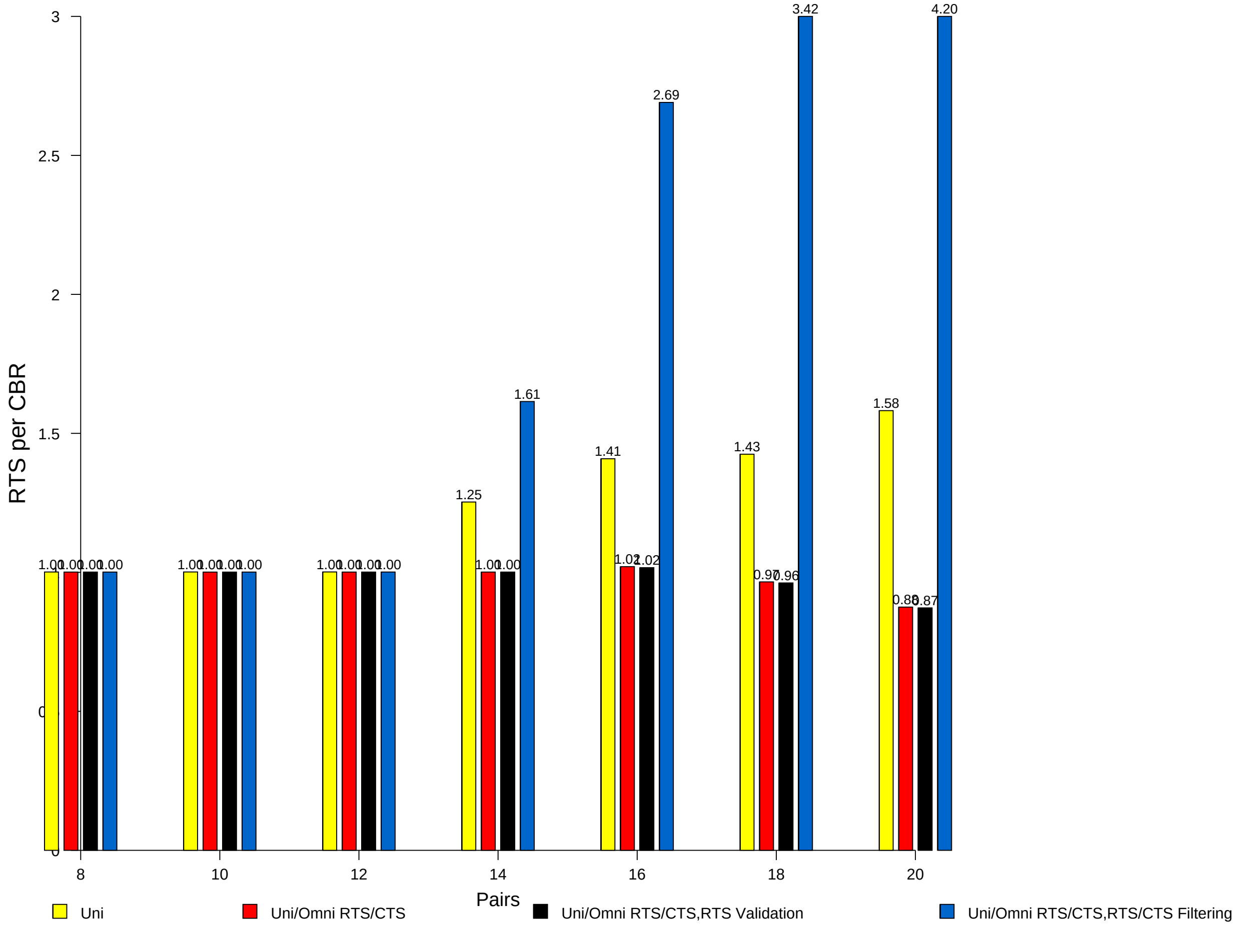

Figure 23: Omni RTS/CTS,Backbone Traffic,RTS per CBR Packet Offered

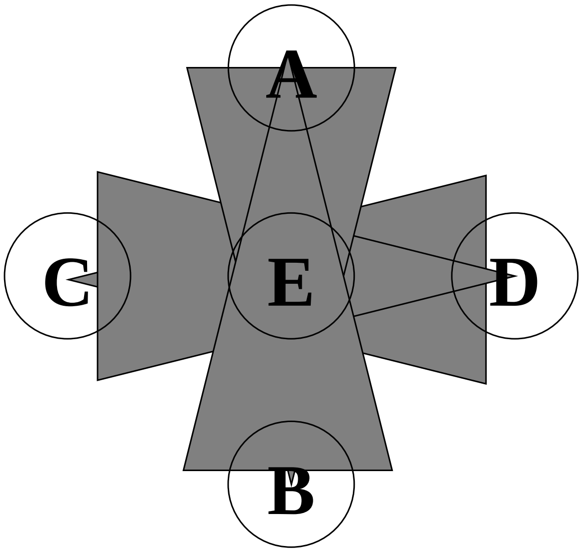

Figure 24: RTS/CTS Filtering Vulnerability at Node E

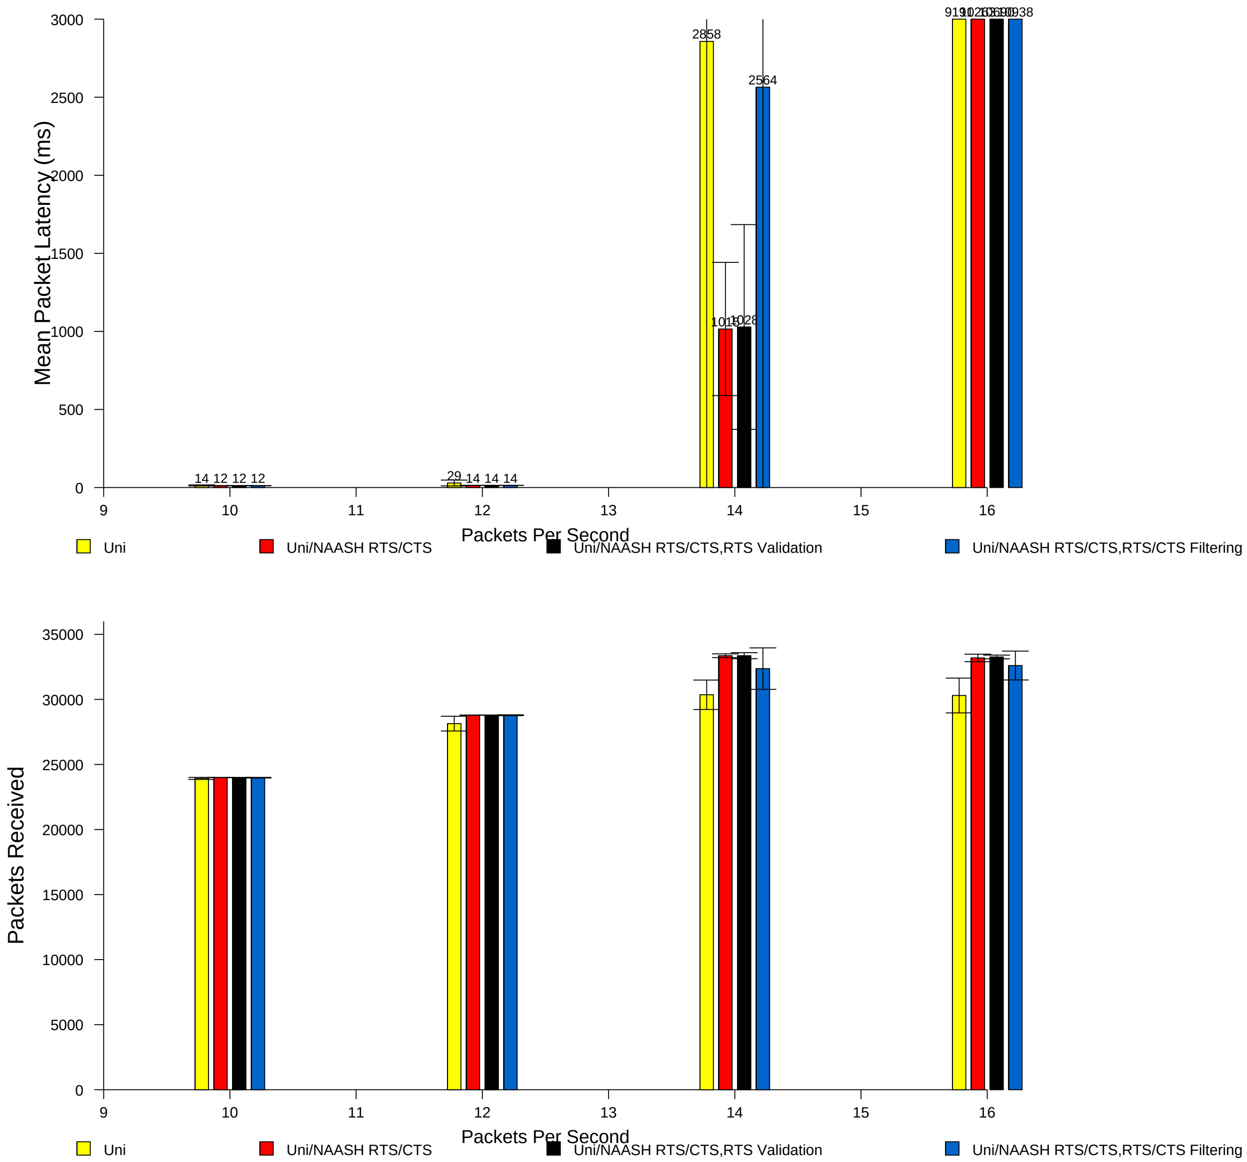

Figure 25: NAASH RTS/CTS,Single Cluster Traffic,Goodput and Latency

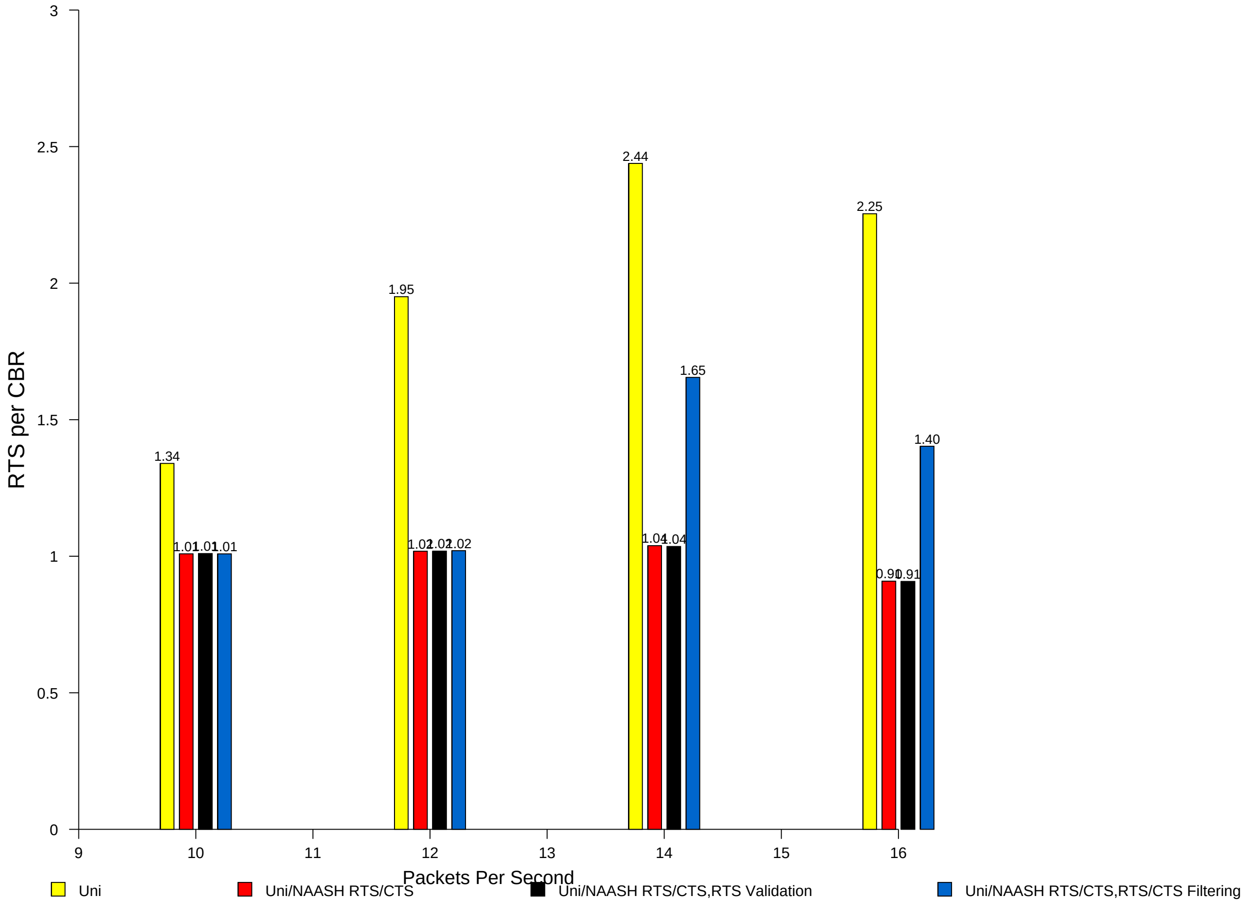

Figure 26: NAASH RTS/CTS,Single Cluster Traffic,RTS per CBR Packet Offered

Since there is only one active cluster in the system, NAASH converges on a gain pattern which is effectively equivalent to omnidirectional. The backbone scenarios exhibit behavior like the simple unidirectional case. This is as expected since simple point-to-point traffic results in the NAASH Active Traffic Table containing only one other node, *i.e.* the destination, which results in the use of purely unidirectional gain patterns. These results are shown in Figures 29 and 30. Utilizing NAASH results in scalability improvements compared with the omnidirectional profile when used in the multiple cluster scenarios. This is illustrated in Figure 29. When used in conjuction with RTS/CTS Filtering it exhibits scalability competitive with that of the simple unidirectional case.

4.6 Overall Results

In general, the best overall performance results occur when stations are either using omnidirectional RTS/CTS (single cluster) or a purely unidirectional profile (backbone and multiple cluster). A very valid question is whether it makes sense to utilize the “in between” patterns, or if as good or better results would be obtained by utilizing a much simpler heuristic. A simplified vesion of NAASH could switch between purely unidirectional and omnidirectional RTS/CTS based simply on the presence of *any* other stations transmitting to the destination or to the source station itself without calculating the specific gain pattern required to cover them. While such a scheme has definite advantages in terms of simplicity and the amount of location information required, it will result in a larger overall RF footprint for stations in many cases, and potentially cause interference and noise problems over a larger area.

5 Conclusions and Future Work

5.1 Conclusions

In this work we have outlined a category of unidirectional enhancements to the 802.11 MAC which maintain interoperability with unmodified omnidirectional equipment. We have also explored the potential benefits of using non-unidirectional antenna profiles in order to decrease deafness while still maintaining a degree of spatial reuse competitive with purely unidirectional profiles. Furthermore, we have evaluated two schemes for mitigating the effect of excessive RTS/CTS messages which arise when using directional antennas.

5.2 Future Work

There are a number of possible avenues to explore in future work. The simple adaptive unidirectional/omnidirectional scheme outlined in §4.6 and its effect on overall spatial reuse compared with a NAASH gain pattern is promising. Another interesting prospect is to leave the antenna in a narrower gain profile while waiting idle for transmission. This corresponds to the second half of Table 2, which requires some foreknowledge that a transmission is going to occur. Leaving the antenna in NAASH configuration while idle could be effective at reducing interference from unrelated traffic while not shutting out related traffic. RTS/CTS Filtering itself can be further refined and tuned. Its propensity for ignoring too many RTS/CTS messages as the traffic load increases could be addressed by modifying the scheme to weigh successful exchanges more heavily than failed ones or by increasing the saturation point. Investigating the effectiveness of a DNAV is also a worthwhile direction. Even though it requires location information about *all* of the surrounding stations to function (the schemes proposed in this work only require location information about neighbors which form part of the same

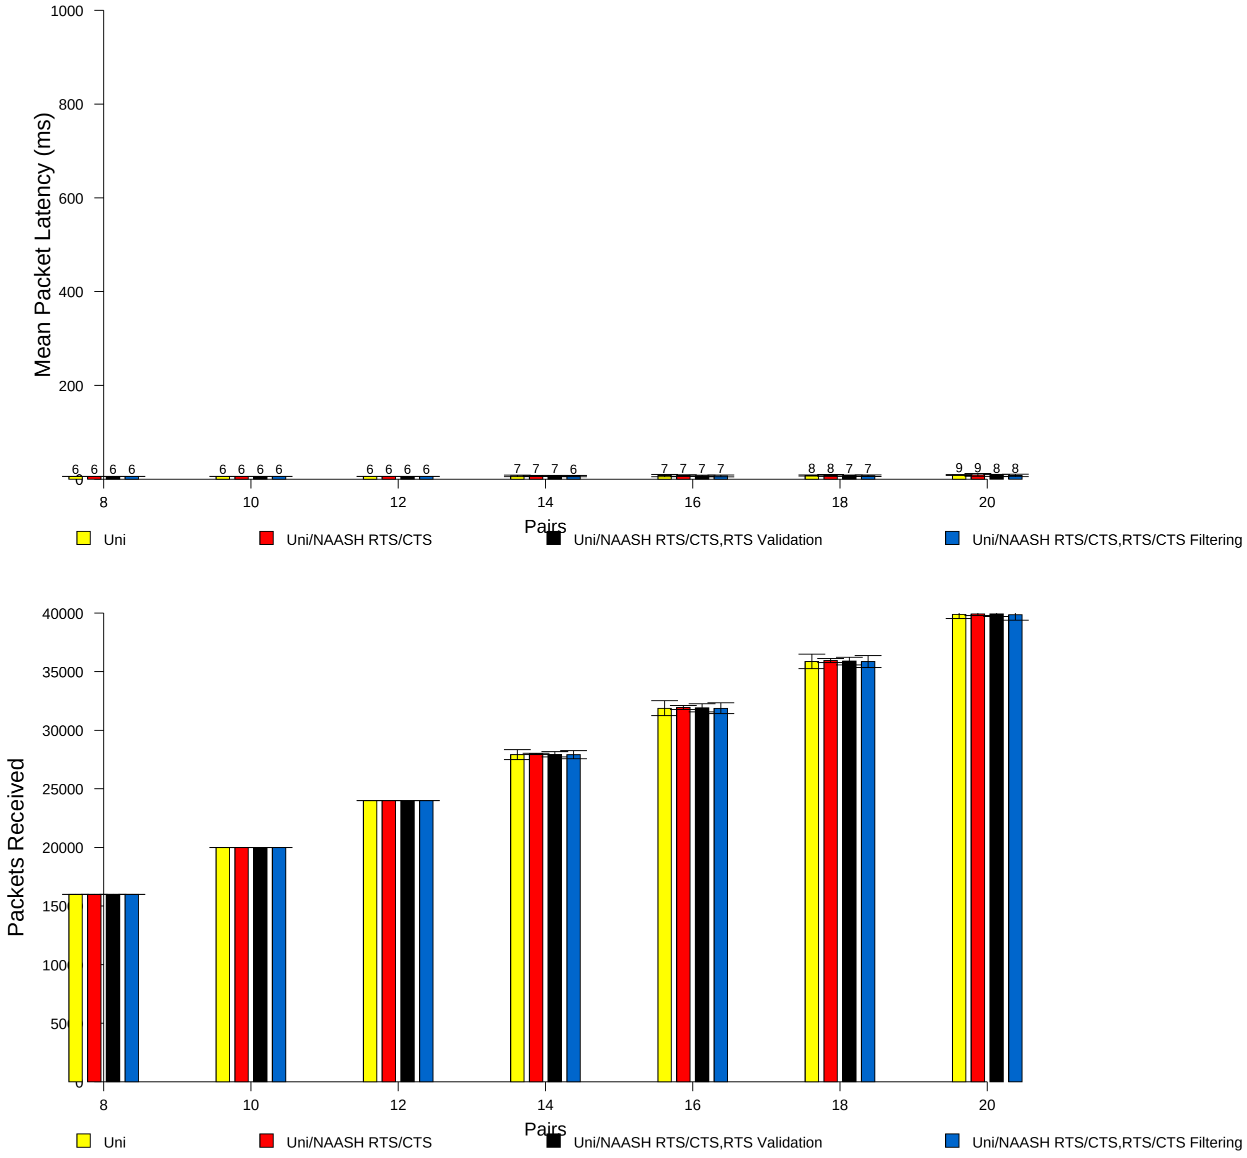

Figure 27: NAASH RTS/CTS,Backbone Traffic,Goodput and Latency

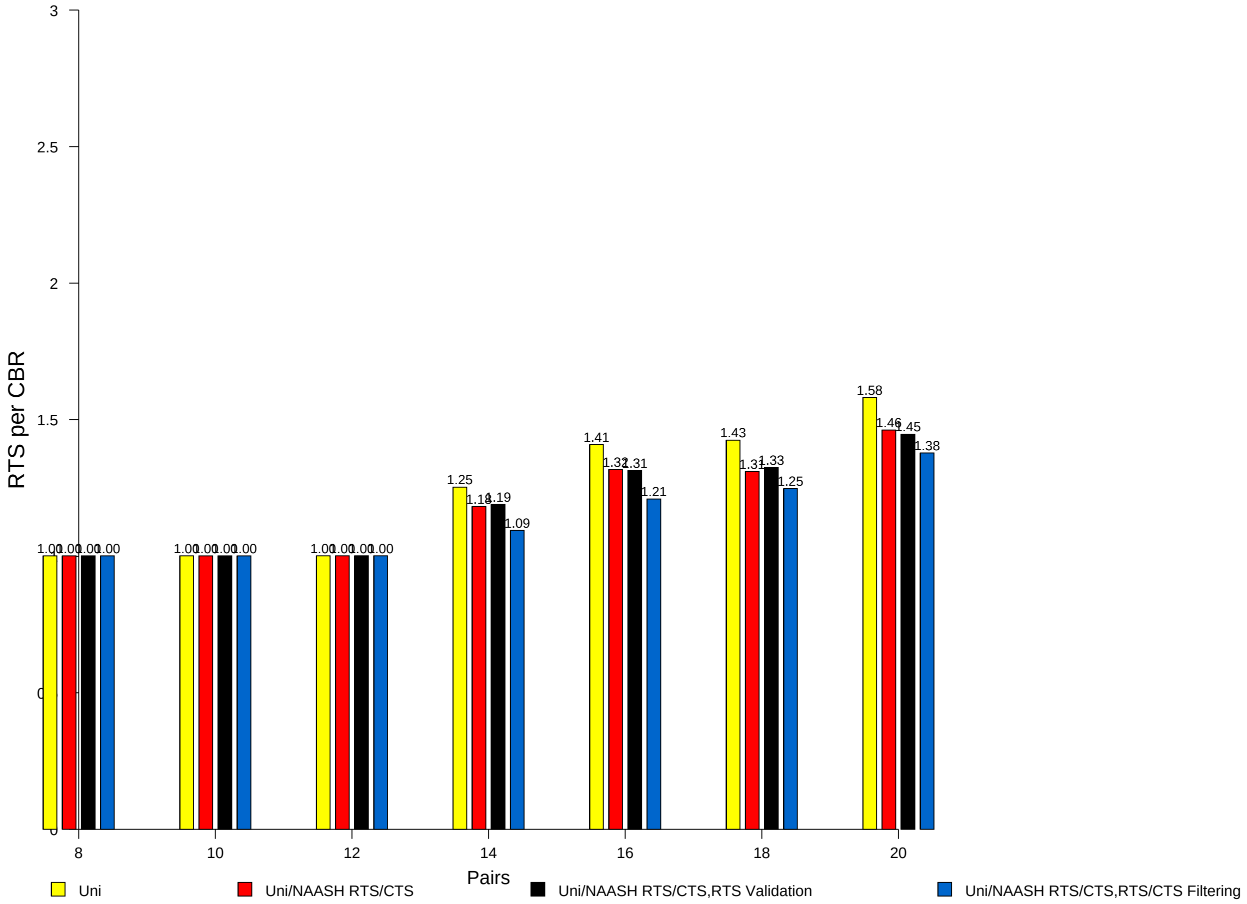

Figure 28: NAASH RTS/CTS,Backbone Traffic,RTS per CBR Packet Offered

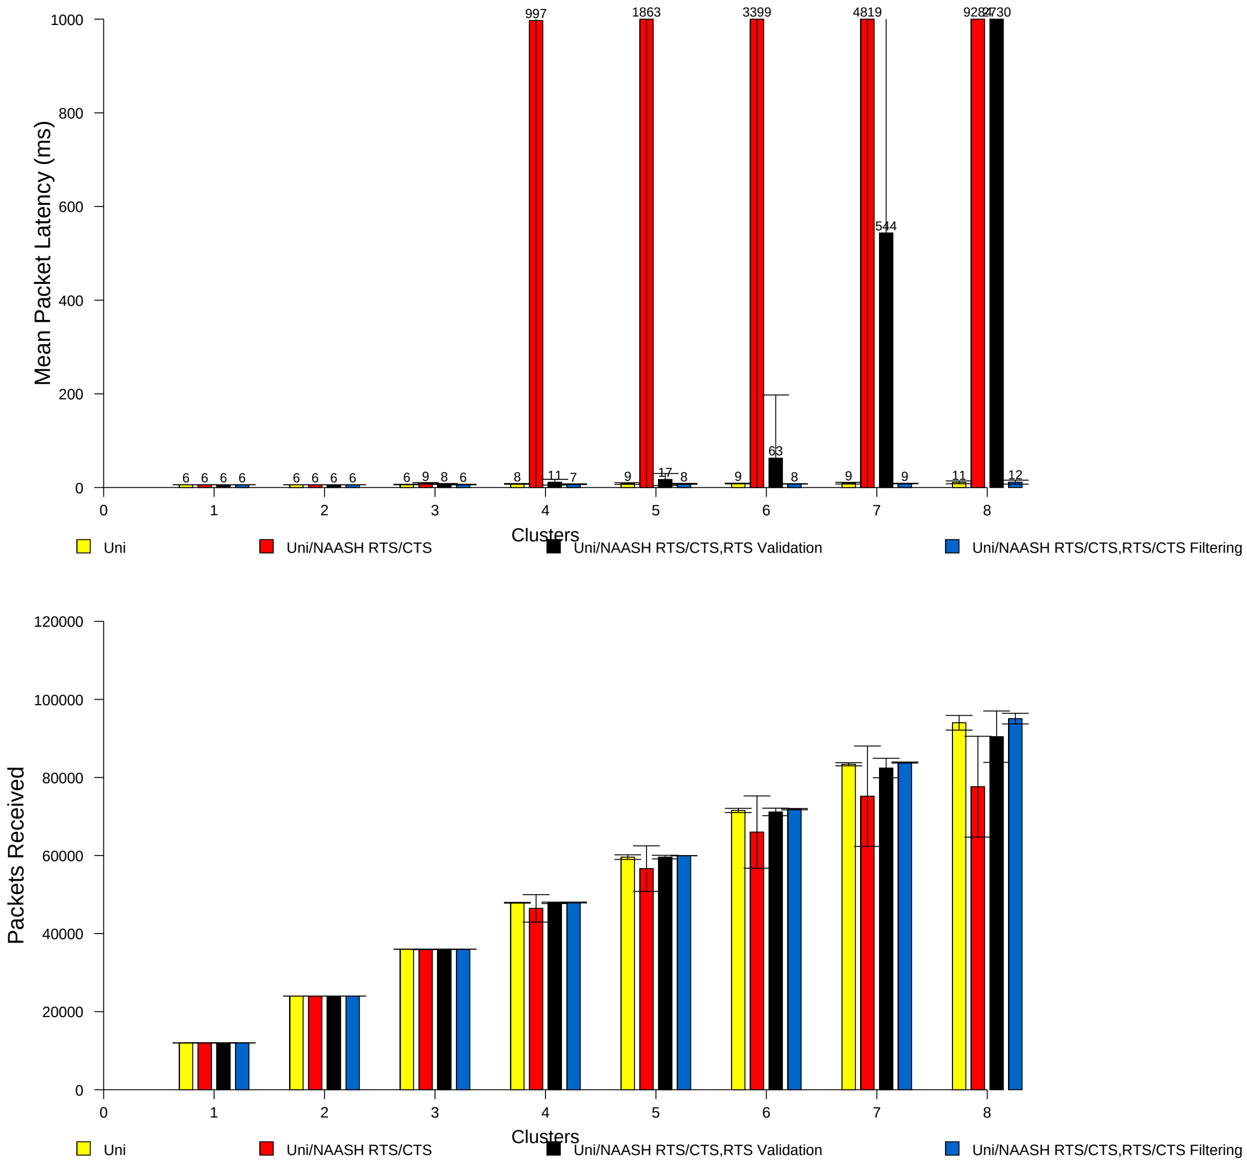

Figure 29: NAASH RTS/CTS,Multiple Cluster Traffic,Goodput and Latency

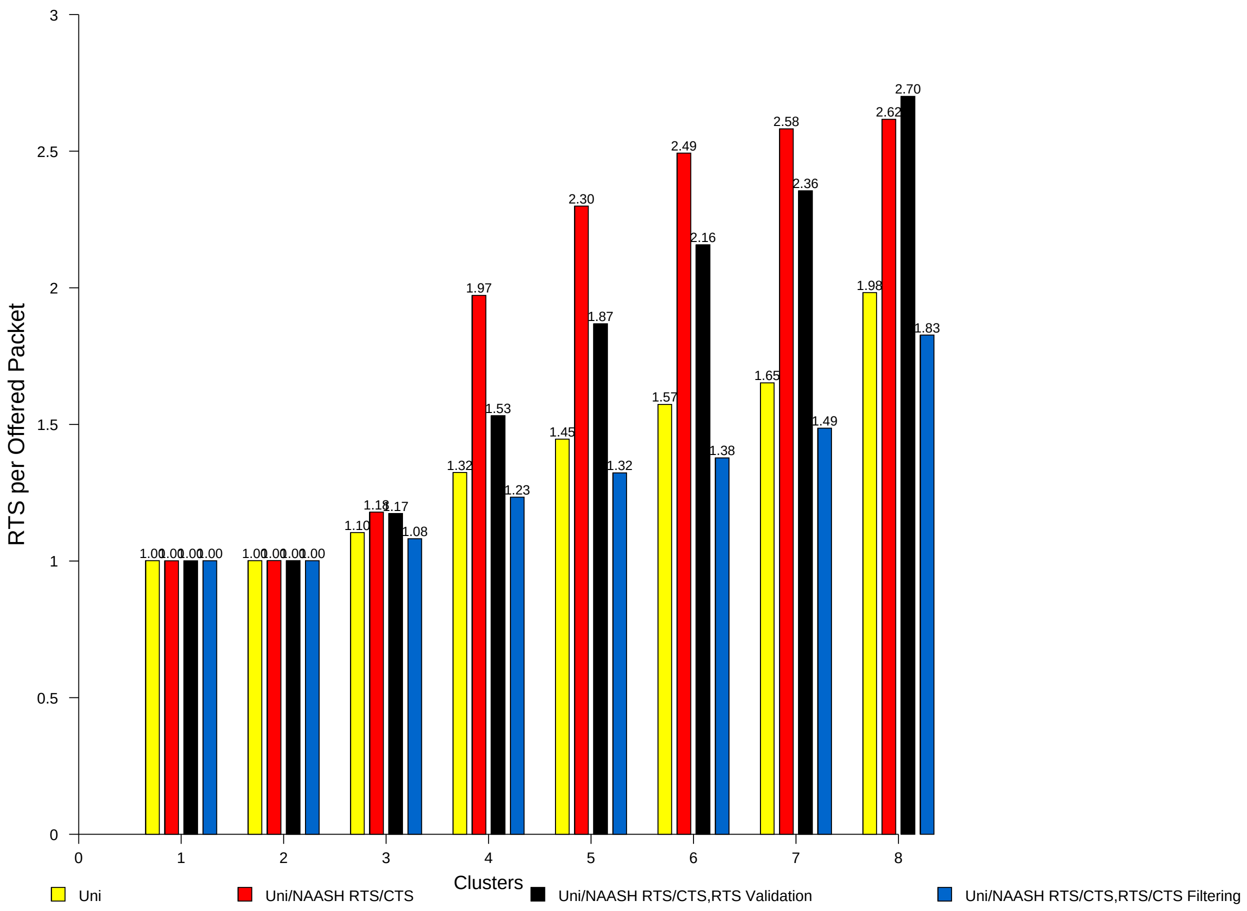

Figure 30: NAASH RTS/CTS,Multiple Cluster Traffic,RTS per CBR Packet Offered

communicating “clique”), it is intended to enhance spatial reuse and could help mitigate the effects of excessive RTS/CTS traffic.

6 Acknowledgements

Computer time was provided by equipment purchased under NSF ARI Grant #CDA-9601817 and NSF sponsorship of the National Center for Atmospheric Research. The authors would also like to thank to Anmol Sheth for providing valuable discussion and feedback.

References

[1] Romit Roy Choudhury and Nitin Vaidya. Deafness: A MAC Layer Problem in Ad Hoc Networks when Using Directional Antennas. Technical report, University of Illinois at Urbana-Champaign, July 2003.

[2] J. Li, C. Blake, D. DeCouto, H.I. Lee, and R. Morris. Capacity of ad hoc wireless networks. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2001)*, Rome, Italy, July 2001.

[3] Saikat Ray, Jeffrey Carruthers, and David Starobinski. Rts/cts-induced congestion in ad-hoc wireless lans. In *Proceedings of the 2003 IEEE Wireless Communications and Networking Conference*, pages 1516–1521. IEEE WCNC, March 2003.

[4] Thanasis Korakis, Gentian Jakllari, and Leandros Tassiulas. A mac protocol for full exploitation of directional antennas in ad-hoc wireless networks. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking and computing*, pages 98–107. ACM Press, 2003.

[5] Siuli Roy, Dola Saha, S. Bandyopadhyay, Tetsuro Ueda, and Shinsuke Tanaka. A network-aware mac and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna. In *Proceedings of the 4th ACM international symposium on Mobile ad hoc networking and computing*, pages 88–97. ACM Press, 2003.

[6] Young-Bae Ko, Vinaychandra Shankarkumar, and Nitin H. Vaidya. Medium access control protocols using directional antennas in ad hoc networks. In *INFOCOM (1)*, pages 13–21, 2000.

[7] Mineo Takai, Jay Martin, Rajive Bagrodia, and Aifeng Ren. Directional virtual carrier sensing for directional antennas in mobile ad hoc networks. In *Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking and computing*, pages 183–193. ACM Press, 2002.

[8] Ram Ramanathan. On the performance of ad hoc networks with beamforming antennas. In *Proceedings of the 2nd ACM international symposium on Mobile ad hoc networking and computing*, pages 95–105. ACM Press, 2001.

[9] Marvin Snchez, Tim Giles, and Jens Zander. Csma/ca with beam forming antennas in multi-hop packet radio.

[10] Yu Wang and J.J. Garcia-Luna-Aceves. Spatial reuse and collision avoidance in ad hoc networks with directional antennas.

[11] Yu Wang and J.J. Garcia-Luna-Aceves. Collision avoidance in single-channel ad hoc networks using directional antennas.

[12] J. Li, J. Jannotti, D. De Couto, D. Karger, and R. Morris. A scalable location service for geographic ad-hoc routing. In *Proceedings of the 6th ACM International Conference on Mobile Computing and Networking (MobiCom '00)*, pages 120–130, August 2000.

[13] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M. Frans Kaashoek. The click modular router. *ACM Transactions on Computer Systems*, 18(3):263–297, August 2000.

[14] Michael Neufeld, Ashish Jain, and Dirk Grunwald. Nsclick: Bridging Network Simulation and Deployment. In *Proceedings of the Fifth ACM Internation Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, 2002.

[15] Microsoft terraserver. available from `http://terraserver-usa.com`.

[16] Baruch Awerbuch, David Holmer, and Herbert Rubens. High throughput route selection in multi-rate ad hoc wireless networks.