

A Configurable and Extensible Transport Protocol

Patrick G. Bridges, *Member, IEEE*, Gary T. Wong, Matti Hiltunen, *Member, IEEE*,
Richard D. Schlichting, *Fellow, IEEE*, and Matthew J. Barrick, *Student Member, IEEE*

Abstract—The ability to configure transport protocols from collections of smaller software modules allows the characteristics of the protocol to be customized for a specific application or network technology. This paper describes a configurable transport protocol system called CTP in which microprotocols implementing individual attributes of transport can be combined into a composite protocol that realizes the desired overall functionality. In addition to describing the overall architecture of CTP and its microprotocols, this paper also presents experiments on both local area and wide area platforms that illustrate the flexibility of CTP and how its ability to match more closely application needs can result in better application performance. The prototype implementation of CTP has been built using the C version of the Cactus microprotocol composition framework running on Linux.

Index Terms—Configuration, customization, extensibility, transport protocol.

I. INTRODUCTION

EXISTING network transport protocols such as TCP and UDP have limitations when they are used in new application domains and for new network technologies. For example, multimedia applications sharing a network need congestion control but not necessarily ordered reliable delivery, a combination implemented by neither TCP nor UDP. Similarly, the congestion control mechanisms in TCP work well in wired networks but often over-react in wireless networks where packets can be lost due to factors other than congestion. The lack of appropriate guarantees or specific features has led to the widespread development of specialized protocols used in conjunction with or instead of standard transport protocols. These include IPsec [1] and SSL [2] for security, RSVP [3] for bandwidth reservation, RTP [4] for real-time audio and video, GTP [5] and CEP [6] for transport in Grid and high-end computing environments, and SCTP [7] for enhanced transport reliability. Developing such a protocol from scratch is, needless to say, often a significant undertaking.

In this paper, we describe our experience building a configurable transport protocol, CTP, that allows protocol semantics to be tuned to specific application needs without the engineering

effort involved with new protocol development. With this approach, software modules that implement different service attributes or variants are written, and then a custom protocol is constructed by selecting appropriate modules based on the needs of the higher levels that use the service or on the specific characteristics of the underlying network or computing platform. Thus, for example, a congestion-control module can be configured together with a datagram service. The net result is, in effect, a family of transport protocols, each useful in a given scenario.

We experimentally demonstrate that CTP achieves comparable performance to existing protocols such as TCP and UDP on the applications for which they were designed. More importantly, we show that CTP can be customized for new applications to provide better performance than existing protocols without the software engineering overhead associated with developing a new protocol from scratch. Our prototype version of CTP is implemented using the Cactus microprotocol composition framework [8] running on UNIX UDP sockets on a cluster of Linux x86 machines and between Linux x86 machines across the Internet.

II. CTP DESIGN

A. Transport Attributes and Algorithms

The first step in developing a customizable transport protocol is identifying various quality attributes and algorithms. These include:

- *Reliability*. Addresses the likelihood that the receiver receives all the data sent by the sender.
- *Ordering*. Describes guarantees concerning the ordering of data at the receiver relative to the order in which it is sent.
- *Performance*. Describes whether data is transported from sender to receiver best effort or with some guaranteed performance using resource reservation.
- *Timeliness*. Describes the timing characteristics of the end-to-end transmission with respect to maximum latency or jitter.

TCP and UDP provide essentially a fixed set of these attributes. In particular, TCP provides strong reliability (guaranteed delivery) and ordering (in-order byte stream) semantics, but only best-effort performance and no timeliness guarantees. Similarly, UDP provides best-effort performance, but with no ordering, timeliness, or reliability guarantees.

Given an attribute, numerous algorithms and protocols are often available for implementing its properties. For example, reliability can use some combination of positive, negative, or selective acknowledgment protocols, or several different forward error correction schemes. In some cases, different algorithms provide different types of guarantees. For example, IP-style one's complement and cyclic redundancy checks

Manuscript received February 12, 2004; revised August 2, 2006, January 30, 2007, and March 29, 2007; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor K. Calvert. This work was supported in part by DARPA under Grant N66001-97-C-8518, NSF under Grants CDA-9500991 and ANI-9979438, DOE Office of Science Grant DE-FG02-05ER25662, and the Sandia University Research Program Contract 190576.

P. G. Bridges and M. J. Barrick are with the Department of Computer Science, University of New Mexico, Albuquerque, NM 87131 USA (e-mail: bridges@cs.unm.edu; patrick.bridges@mac.com; barrick@cs.unm.edu).

G. T. Wong is with the Department of Computer Science, Boston University, Boston, MA 02215 USA (e-mail: gtw@cs.bu.edu).

M. Hiltunen and R. D. Schlichting are with the AT&T Labs—Research, Florham Park, NJ 07932 USA (e-mail: hiltunen@research.att.com; rick@research.att.com).

Digital Object Identifier 10.1109/TNET.2007.906245

(CRC) provide integrity that protects against accidental data modification, while cryptographic methods protect against intentional modification. In other cases, different algorithms provide approximately the same guarantee, but with different trade-offs with respect to resource usage or other attributes. For example, forward error correction uses more bandwidth than acknowledgments, but provides faster recovery from failures, and thus a smoother data flow at the receiver.

Different choices can also be made for other design elements, such as whether to use congestion and/or flow control, and if so, what type. The protocol must also be able to interact appropriately with the protocol below it in the protocol stack. For example, messages may need to be fragmented into pieces or small messages coalesced into one packet. If a resource reservation protocol such as RSVP is available, the transport protocol may interact with it to make a resource reservation for the connection. Finally, the transport protocol must deal with such practical issues as connection establishment, monitoring, and tear-down.

B. Cactus

Cactus is a system for constructing highly-configurable protocols for networked and distributed systems [8]. Individual protocols in Cactus, termed *composite protocols*, are constructed from fine-grained software modules called microprotocols that interact using an event-driven execution paradigm. Each microprotocol is structured as a collection of event handlers and generally implements a distinct property or function of the protocol. Composite protocols are then layered on top of each other to create a protocol stack using an interface similar to the standard *x*-kernel API [9]. This two-level approach has a high degree of flexibility, yet provides enough structure and control that it is easy to build collections of modules realizing a large number of diverse properties.

At runtime, composite protocol instances, termed *composite sessions*, are used to process packets. Composite sessions are created by protocol routines in response to open requests from either local applications or received packets. Each composite session contains a collection of *microprotocol instances* in which event handlers are bound to protocol-specific events to effect protocol processing.

The programming model in Cactus is based on events and event handlers. Events are used to signify state changes of interest, such as “message arrival from the network”. When such an event occurs, all event handlers bound to that event are executed. Events can be raised explicitly by microprotocol instances or implicitly by the runtime system. The runtime system also provides a variety of operations for managing events and event handlers. In addition to traditional blocking events, Cactus events can also be raised with a specified delay to implement time-driven execution, and can be raised asynchronously. Arguments can be passed to handlers in two ways, statically when a handler is bound to an event and dynamically when an event is raised. Other operations are available for unbinding handlers, creating and deleting events, halting event execution, and canceling a delayed event. Handler execution is atomic with respect to concurrency, i.e., a handler is executed to completion before any other handler is started unless it voluntarily yields the CPU.

The Cactus message abstraction is designed to facilitate development of configurable services. One of the main features of Cactus messages are *message attributes*, which are a generalization of traditional message headers. Operations are provided for microprotocols to add, read, and delete message attributes. Furthermore, a customizable pack routine combines message attributes with the message body for network transmission (on-wire format), while an analogous unpack routine extracts attributes at the receiver.

Synchronization and coordination of execution activities in Cactus is accomplished through *event-based barriers* that can be associated with data items, including messages. A microprotocol instance can register with the barrier, and an event associated with the barrier will only be raised when all microprotocol instances registered with the barrier have entered the barrier. These barriers are used to coordinate activities across multiple microprotocols, especially to control the transfer of messages up and down the protocol stack.

C. Design Overview

CTP is a composite protocol in which each attribute or function described in Section II-A is implemented by one microprotocol or a set of alternative microprotocols. Thus, the current design has one or more microprotocols for reliability, ordering, security, jitter control, congestion control, flow control, data and header compression, MTU discovery, message fragmentation and collation, and connection establishment, monitoring and tear-down. The goal of the design is to decouple the implementations of different attributes and functions to maximize the ability to mix and match different microprotocols to provide exactly the required properties. Decoupling the different features of transport protocols is not trivial, since often much of the functionality is tightly coupled for efficiency. For example, reliability, congestion control, and flow control in TCP often use the same transmission window data structure, while byte sequence numbers are used to implement reliability, ordering, and flow control feedback.

The current CTP design focuses on only bidirectional message-oriented point-to-point communication over an unreliable packet-oriented network protocol (e.g., IP). Specifically, an application uses a given CTP configuration to exchange arbitrary length messages (e.g., a video frame) with some application-defined semantics with a single endpoint. Since the design of CTP does not assume that the underlying network protocol supports such arbitrary length messages, microprotocols for fragmenting or coalesces messages into an appropriate transport unit—a *segment*—are provided. Finally, CTP addresses are currently local/remote IP/port number 4-tuples similar to those used by TCP.

D. CTP Events

Microprotocol instances in a CTP session interact using events in order to manipulate shared data, largely messages and their attributes. Fig. 1 shows the predefined set of common events useable by all CTP microprotocols; solid arrows are used to indicate events raised by CTP’s interface routines and dashed arrows to indicate causal relations between other events. For example, when the MSG FROM USER event is raised,

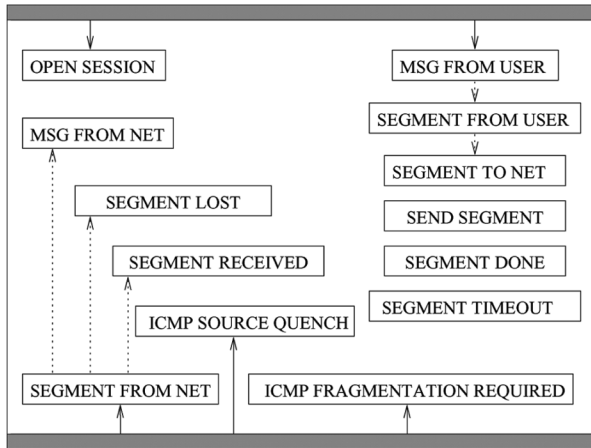


Fig. 1. Major CTP events.

some microprotocol will raise the `SEGMENT FROM USER` event. Additional local timeout events are used by several of the microprotocols. Most of the event names are self explanatory; for example, `SEGMENT RECEIVED` is raised when a previously sent segment is acknowledged.

CTP makes extensive use of the Cactus event-based barrier, particularly event-based barriers associated with individual messages. For historical reasons, event-based barriers associated with messages are generally referred to as *hold bits*. CTP uses three sets of hold bits on each message: send bits, deallocate bits, and done bits. The send bits are used to coordinate sending of segments down to lower layers and delivery of messages up to the application. A microprotocol sets a send bit in a given message when the message can exit the composite protocol as far as the microprotocol is concerned. When all bits in a message are set, it exits and the `SEND SEGMENT` event is raised to notify microprotocols that the segment has actually left the protocol. For example, congestion control, flow control, and reliability functionality in CTP each control send bits to determine when a segment can be transmitted.

Send bits allow different microprotocols to operate on messages independently without knowing which other microprotocols need to process the message. They also decouple the approval process from any kind of ordering—when all the required microprotocols have set their bits, the message exits the composite protocol independent of the order in which they were set. Note that systems supporting only hierarchical composition intrinsically dictate one fixed release order. Similarly, deallocate bits are used for determining when a segment will not be needed by any microprotocol and thus can be deleted.

Done bits are used when agreement from multiple microprotocols is needed on a property unrelated to a message either exiting the composite protocol or being deallocated. For example, a congestion control microprotocol needs to know when an outgoing message is not on the network (i.e., has either been acknowledged or timed out) *and* will never be retransmitted in order to advance the trailing edge of the congestion control window. To do this, every relevant microprotocol sets a done bit at the appropriate time, and then the `SEGMENT DONE` event is raised when all bits are set. Note that deallocate bits are not suf-

ficient for this purpose because microprotocols may delay setting deallocate bits on messages even though they will never retransmit the packet. For example, the Forward Error Correction microprotocol delays deallocating sent messages so that it has the data needed to compute the contents of redundant packets using an erasure code algorithm.

E. Configuration and Initialization

CTP composite sessions are created in response to an explicit open request from an application or when a packet with a host/port 4-tuple is received that does not demultiplex to an existing session. The CTP session initialization routine is then invoked, resulting in the creation of session-global state, the instantiation of microprotocol instances, and the initialization of these microprotocol instances. At this time, microprotocol instance initialization routines set up their data structures and notify the runtime system of any necessary hold bits they will need on CTP messages. After the session and all of its microprotocol instances are initialized, the CTP demultiplexing routine raises the `OPEN SESSION` event in the new session so that microprotocols that perform connection establishment can execute appropriately. If the session was created as a reaction to a packet received from the network, the `SEGMENT FROM NET` event will be raised to allow processing of any data contained in the packet.

New CTP sessions select the appropriate microprotocol instances for each composite session based either on information in the locally-generated open request or on data in the packet that caused the creation of the new session. For local open requests, the current CTP implementation requires applications to specify exactly the microprotocols they desire in the session being created, including resolving dependencies by hand. Configuration tools such as those used in previous Cactus-based systems [10] could also be used to ease this process.

For open requests received from a remote host, CTP requires that packets that create a session contain sufficient data to determine which microprotocols were used to generate the received segment. In the most general case, connectionless protocols where any packet can establish a session, this is implemented as a 32-bit bitfield that is included with every CTP packet, with a different bit assigned to each possible CTP microprotocol. For connection-oriented CTP configurations, however, this bitfield need only be included in the connection establishment request.

Note that the microprotocol configuration in an existing CTP session is not currently changeable at runtime. While feasible in principle, doing so would require substantial additional machinery either to quiesce the network or to support multiple microprotocol instances simultaneously while old packets are drained from the network. However, applicable research into supporting such dynamic adaptation capabilities has been done both in the context of Cactus [11], [12] and in other systems such as the K42 operating system [13].

III. MICROPROTOCOL HIGHLIGHTS

A. Base Functionality

Transport Driver is the only microprotocol that must be present in any configuration. It adds port identifiers on all

outgoing segments for demultiplexing and also contains trivial handlers for certain events to ensure that a message is carried through CTP irrespective of the presence of other microprotocols.

The Sequenced Messages and Sequenced Segments microprotocols add message attributes uniquely identifying each outgoing message and segment, respectively. While this labeling does not provide any service to the application, it is useful for other microprotocols such as reliability or ordering. Performing the procedure in a separate microprotocol allows the other microprotocols to share the same attribute, saving space in the message.

A group of microprotocols transforms messages into segments at the sender and then back to messages at the receiver. They are also responsible for raising the SEGMENT FROM USER and MSG FROM NET events. Fixed Size simply creates a separate segment from each message, Coalesce combines multiple small messages into one segment, and Resize fragments the messages into segments that can be handled by the underlying IP network without IP-level fragmentation (MTU discovery). One of these microprotocols must be present in each configuration.

Finally, a set of optional microprotocols is responsible for establishing and shutting down a connection, and for monitoring its status. Virtual Circuit implements a handshake protocol that provides reliable startup and shutdown semantics, and exchanges random initial sequence numbers for message and segment numbering. Virtual Circuit is completely transparent to other microprotocols, even those that use sequence numbers—if it is not included, constant initial values are used. An additional Keep Alive microprotocol is responsible for sending probe messages to detect link failures in the absence of application messages.

B. Informational Microprotocols

CTP contains a number of microprotocols that collect information and provide it to other microprotocols by raising events or setting shared variables. For example, the Round Trip Time Estimation microprotocol maintains an estimate of the end-to-end round trip time in a protocol-wide shared variable by handling the SEND SEGMENT and SEGMENT FROM NET events so that it can note when a segment is actually placed on the wire and when an acknowledgment for the segment is received. This estimate is then used by other microprotocols for detecting congestion and setting timeout values, for example.

The Positive ACK microprotocol is another, more complex, informational microprotocol used to track the status of transmitted segments. It implements a general cumulative acknowledgment facility necessarily more general than similar functionality in other protocols. In particular, it can be used in CTP configurations that do not provide reliable delivery guarantees because it does not incorporate functionality such as retransmissions. This allows it to be used, for example, in unreliable protocols that still need to track packet delivery status for flow and congestion control purposes, as well as in reliable configurations that include microprotocols such as Retransmit.

The generality of Positive ACK is achieved by slightly redefining the meaning of a cumulative acknowledgment and in-

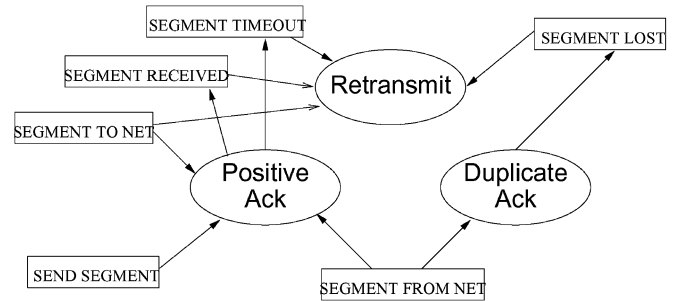


Fig. 2. ACK-related event handling.

roducing a session-global data structure to decouple the microprotocol from the presence of reliability microprotocols. In Positive ACK, an acknowledgment indicates that the acknowledged segment was received *and* that the receiver no longer needs or expects to receive the acknowledged segment or any segment sent prior to it. Note that this does not mean that the previous segments were necessarily received—simply that they are unneeded, that is, that their *reliability constraints* have been met.

The session-global data structure used by Positive ACK keeps track of whether the reliability constraints on each packet have been met. If a reliability microprotocol is included in CTP, it sets the default reliability status of packets to RELIABILITY_UNMET in its initialization routine, and then later sets it to RELIABILITY_MET when the packet is acknowledged. If a reliability microprotocol is not included in the configuration, however, the default status of packets remains RELIABILITY_MET. This allows Positive ACK and similar informational microprotocols to know for which packet to send a cumulative acknowledgment.

In reliable protocols where the receiver expects to receive every packet, the more general definition of acknowledgments and the reliability tracking data structure results in the standard acknowledgment behavior used in protocols such as TCP. In protocols that do not require complete reliability, however, the more general definition allows acknowledgments for packets to be sent even if some previous packets have not been received. In addition, this design also allows for partially reliable configurations where some packets must be transported reliably and some unreliably, although CTP does not currently include any microprotocols that make use of this flexibility.

Fig. 2 shows how events are used to track segment status in acknowledgment processing; an arrow pointing to a microprotocol indicates that it has a handler bound to the event, while an arrow originating at a microprotocol indicates that it raises the event. For each outgoing message (event SEGMENT TO NET), Positive ACK includes a cumulative acknowledgment attribute as described above, and also raises the SEGMENT TIMEOUT timer event when the message is actually transmitted (event SEND SEGMENT). For each incoming message (event SEGMENT FROM NET), it checks the acknowledgment attribute, and cancels the SEGMENT TIMEOUT event and raises the SEGMENT RECEIVED event if appropriate. Similarly, the Duplicate ACK microprotocol also monitors the SEGMENT FROM NET event and raises the

```

micro-protocol Retransmit () {
  handler handleSegToNet(segment s){
    setSendBit (s);
  }

  handler handleRetransmit(int seq){
    m = HashLookup(protocolState.segmentHash, seq);
    clearSendBit (m); setSendBit (m);
  }

  handler handleSegRxd(int seq){
    m = HashLookup(protocolState.segmentHash, seq);
    SetReliabilityState (protocolState.relList, seq,
      RELIABILITY_MET);
    setDoneBit (m); setDeallocateBit (m);
  }
}
initial {
  protocolState.defaultRelStatus = RELIABILITY_UNMET;
  requestDoneBit ();
  requestDeallocateBit ();
  requestSendBit ();
  bind(SEGMENT TO NET,handleSegToNet,0);
  bind(SEGMENT LOST,handleRetransmit,0);
  bind(SEGMENT TIMEOUT,handleRetransmit,0);
  bind(SEGMENT RECEIVED,handleSegRxd,0);
}
}

```

Fig. 3. Retransmit microprotocol pseudo code.

SEGMENT LOST event when appropriate. Other microprotocols such as Retransmit then use these events and the data included in event arguments (e.g., sequence numbers) to determine when to retransmit old segments or release new packets to the network.

C. Reliability Microprotocols

Reliable transmission can be implemented using different types of redundancy ranging from redundant network connections to redundant transmission over the same connection. CTP currently has two reliability microprotocols: Retransmit and Forward Error Correction. Retransmit is a traditional ARQ reliability scheme that relies on informational microprotocols to know when packets have been received and which ones should be retransmitted. As shown in the pseudocode in Fig. 3, it handles the SEGMENT LOST and SEGMENT TIMEOUT events and retransmits the appropriate segment when one of these events is raised. In addition, it allocates a done bit on each outgoing message and sets it upon receiving the SEGMENT RECEIVED event. As mentioned in Section II, this allows other microprotocols to know when the message will not be retransmitted so that they can, for example, advance the congestion window.

Forward Error Correction transmits redundant data so that the receivers can reconstruct a complete transmission despite message losses. Forward Error Correction at the receiver then handles the redundant segments and uses them to create a new message for each of these missing segments and raises the SEGMENT FROM NET event for the reconstructed segments. Redundant data packets are also tagged with a special attribute to assure that they are not handed to the application. As a result, other microprotocols see the reconstructed segments as if they had arrived normally. The specific error correction scheme currently used by this microprotocol is a block erasure code algorithm [14] that encodes k segments of original data into n segments of encoded data ($k < n$).

Note that forward error correction and ARQ reliability can often be used together in the same CTP configuration. This gives the user a rich set of possibilities for reliable communication that can be used to match the specific requirements of particular applications.

D. Transmission Control Microprotocols

CTP offers flexible facilities for controlling the speed of transmission, typically used to ensure that a sender limits its outgoing traffic to a level acceptable to the network and receiver. Our architecture divides these microprotocols into two categories: flow control and congestion control.

Flow control refers to end-to-end transmission control that provides a mechanism for the receiver to dictate the sender's transmission speed. Available microprotocols include:

- *XON/XOFF*. The receiver issues suspend/resume instructions to the sender.
- *RTS/CTS*. The sender explicitly requests the ability to send more packets.
- *Windowed*. The receiver periodically informs the sender of its available buffer space.

These microprotocols all operate at the sender side by binding a handler to the SEGMENT TO NET event, which sets its send bit on an outgoing message only when the requirements for transmitting the segment are satisfied.

At the receiver side, there are facilities in the API to allow higher level protocols to specify policies on traffic rates. The flow-control microprotocols can communicate this information to the sender either by transmitting new feedback messages to the sender or by piggybacking the information on existing messages. This feedback is handled at the sender in a handler bound to the SEGMENT FROM NET event.

Congestion control behaves similarly to flow control in that it limits the transmission rate of senders, but is intended to avoid overrunning the capacity of the network rather than the receiver. Congestion control in CTP consists of two types of microprotocols, congestion control microprotocols that implement the mechanisms for controlling congestion and congestion policy microprotocols that describe corresponding policies. Typical configurations would include one congestion control and one congestion policy microprotocol.

Congestion control microprotocols, like those that do flow control, use send bits to regulate segment transmission. These microprotocols monitor protocol-wide shared variables that congestion policy microprotocols change in response to policy-specific indications of congestion. CTP currently implements two congestion control microprotocols: Windowed Congestion Control and Rate-based Congestion Control.

Windowed Congestion Control implements a simple window-based scheme that limits the number of unacknowledged packets in the network. The size of the window is stored in a shared variable that can be changed by congestion policy microprotocols in response to various events. The Rate-based Congestion Control microprotocol works similarly, but instead controls the average outgoing byte rate based on a shared variable. Because each congestion control microprotocol uses a different send bit for controlling segment

transmission, multiple congestion control microprotocols can be used simultaneously when appropriate.

Congestion policy microprotocols work by changing shared variables exported by congestion control microprotocols, and as such, are designed to work with specific congestion control microprotocols. Available policy microprotocols include:

- TCP Congestion Detection. This microprotocol handles the SEGMENT RECEIVED, SEGMENT TIMEOUT, and SEGMENT LOST events and changes the congestion window used by the Windowed Congestion Control in response to these events in accordance with the congestion control policy used by TCP [15]. Note that the policy implemented by this microprotocol does not depend on the presence of the Retransmit microprotocol in the CTP configuration, so it may be used with unreliable communication or in combination with Forward Error Correction.
- TCP-Friendly Rate Control. This microprotocol monitors segments status events and sets the maximum outgoing data rate used by Rate-based Congestion Control according to the TCP response equation [16].
- SCP Congestion Detection. This microprotocol monitors the average round-trip time and the packet status events and sets both the outgoing data rate used by Rate-based Congestion Control and the window size used by Windowed Congestion Control similarly to the SCP protocol [17].

Other congestion policy microprotocols are also easily implemented in this framework.

E. Ordering and Jitter Control Microprotocols

Ordering microprotocols are relatively simple for point-to-point communication as currently supported by CTP. The sender can add a message attribute that indicates the order of the message either as a sequence number or by specifying the message's logical predecessor(s). The current implementation has a Reliable FIFO microprotocol, which enforces strict in-order delivery by buffering out-of-order messages and sending them to the application only after their predecessors have been delivered, and a Lossy FIFO alternative that discards messages that arrive out of order after a configurable delay. A Semantic Order microprotocol uses ordering information provided by the application to record and enforce the logical predecessors of each message. An Out of Band microprotocol can be used with any ordering microprotocol to allow urgent out-of-band messages to be delivered as quickly as possible by overriding the send bit used by the current ordering microprotocol.

Jitter control microprotocols are structurally similar to ordering microprotocols, but use the passage of time rather than predecessor information to decide when the send bit in a message is set. These microprotocols include Fixed Rate Jitter, which delivers messages separated by a fixed time interval and Timestamp Jitter, which preserves the sender's time intervals between messages at the receiver.

IV. DESIGN AND IMPLEMENTATION EXPERIENCES

Over the course of designing and implementing CTP, we gained substantial experience in dealing with configurability in

CTP, as well as using the Cactus protocol framework. We also ran into issues with our original CTP design [18] that had to be resolved. In this section, we discuss these experiences.

A. Configurability and Extensibility in CTP

To make CTP highly configurable, the different microprotocols have been designed to be as independent as possible. However, there are some *dependencies*—when one microprotocol requires that another be in the configuration to function correctly—and some *conflicts*—when two microprotocols cannot be in the same configuration. The dependencies in the current design are relatively simple. Every configuration must have Transport Driver and one of the message-to-segment conversion microprotocols Fixed Size, Coalesce, or Resize. The reliability and FIFO ordering microprotocols use sequence numbers provided by the Sequenced Messages and Sequenced Segments. Similarly, most flow and congestion control microprotocols require an informational microprotocol such as Positive ACK to provide feedback on the status of transmitted segments. Finally, congestion control policy and mechanism microprotocols must be used in conjunction with each other.

Conflicts are either syntactic or semantic in nature. An example of a syntactic conflict is that only one message-to-segment conversion microprotocol should be in each configuration, while an example of a semantic conflict is that Lossy FIFO and a reliable communication microprotocol should not be used together. Semantic conflicts do not cause the combination to fail, but the resulting semantics do not satisfy the properties of both of the microprotocols.

Despite these dependencies and conflicts, there are still hundreds of possible different CTP configurations even with a small number of different microprotocols for each transport property and function. The challenge is to identify the correct configuration for each application domain and execution environment. In many cases, this may require experimentation with different combinations to reach the optimal one.

CTP is also designed to be easily extensible, meaning that new microprotocols can be added without modifying the existing ones. The actual effort needed depends on the type of extension. It is typically trivial to add a new alternative implementation for an existing property or function, since the event and data structure interactions are usually the same as in existing microprotocols.

On the other hand, adding a completely new property or function can be more difficult. The implementor must first determine if CTP already has all the necessary events required by the new microprotocol. If not, the CTP framework or some of the existing microprotocols may need to be modified to raise these events. However, completely new microprotocols can often be implemented using the existing set of events. For example, in our design, the jitter control microprotocols were added after the rest of CTP was designed with no modifications to other microprotocols.

B. Corrected Design Mistakes

Over the course of designing and implementing CTP, we made two substantial design mistakes that required re-archi-

tecting parts of the system. This resulted in an insufficiently flexible protocol when we initially tried to use CTP for multimedia applications such as the one used in the experiments in Section V. These two issues are discussed in more detail below.

Decomposing Complex Interactions: Flexible configurability in CTP did not come without substantial effort. For example, while the reliability, ordering, and flow control transport functions that are tightly connected in TCP are completely independent in our final design, this was not originally the case. In our original design [18], a single Positive ACK microprotocol performed two logically separate functions: tracking the status (*received/lost/timed out*) of transmitted segments, and reliable transmission of segments using timeouts and retransmissions. This overloading, the result of failing to completely decompose acknowledgment functionality inspired by TCP, caused problems for applications that wanted segment status tracking but not retransmissions such as streaming multimedia transmission applications.

Decoupling these responsibilities required the introduction of several new microprotocols and events. Much of this decoupling comes from the use of Cactus' event-based programming model, but some required the generalization of protocol functionality and the introduction of additional mechanisms and data structures. We decomposed the original PositiveACK microprotocol into several microprotocols, namely PositiveACK, DuplicateACK, NegativeACK, Retransmit, and RoundTripTimeEstimation. We also introduced three new events, SEGMENT RECEIVED, SEGMENT LOST, and SEGMENT TIMEOUT, to announce when segments are acknowledged, explicitly lost, or have had an unknown status for an unacceptable amount of time. This decomposition allowed CTP to be configured to use acknowledgments for feedback about segment arrival and loss without mandating the introduction of retransmissions and their negative effects on multimedia applications.

The new Positive ACK microprotocol implements acknowledgments and segment timeouts, while the Negative ACK and Duplicate ACK add additional packet tracking functionality. On the sender side, all of these microprotocols work by raising the appropriate events at the appropriate time; these events are then responded to by Retransmit. On the receiver side, Positive ACK was changed to acknowledge packets when it has either received a packet or no longer needs a packet, and a session-global data structure describes whether the reliability constraints on each received packet have been met. In reliable protocols, where the receiver expects to receive every packet, this behavior results in the standard acknowledgment behavior used in protocols such as TCP. In protocols that do not require complete reliability, however, the more general definition allows an acknowledgment for a packet to be sent even if some previous packets have not been received. Note that this change also required the introduction of done bits for use by WindowedCongestionControl as described in Section III-B.

Separating Mechanism and Policy: Another shortcoming of the original CTP design was that it did not separate congestion control mechanism and policy. As in all systems, keeping such separation is useful. This problem was solved by introducing two different microprotocols that implement congestion control mechanisms, WindowedCongestionControl and

RateBasedCongestionControl, and a variety of different microprotocols that implement different congestion control policies as previously described. The most substantial change required by this generalization was the introduction of the done bits on each CTP segment and the corresponding SEGMENT DONE CTP event, allowing WindowedCongestionControl to advance the trailing edge of the congestion window at the appropriate time.

As a result of this experience, separate policy microprotocols were similarly used for controlling forward error correction parameters when CTP was later modified to support adaptation of error correction parameters. To further enable careful separation of mechanism and policy, later work on a system named Cholla [12] explicitly separated protocol policies into a policy control engine where they could be separately composed, controlled, and analyzed.

C. Cactus Event Experiences

After implementing a variety of CTP microprotocols and testing a variety of different configurations, we found the largest source of bugs was in the ordering of event handlers. Cactus allows event handlers to bind with different *order priorities*, and handlers are run in numeric order priority. Excessive use of event ordering, however, resulted in a number of different bugs. In the original implementation, for example, there were not separate SEGMENT TO NET and SEGMENT SENT events; microprotocols that wanted to run after segments were sent would simply bind to SEGMENT TO NET with a large order priority. As new microprotocols were introduced, however, misorderings between when handlers were run could cause, for example, round trip times to be calculated inappropriately.

To address this problem, later implementations of CTP were changed to use more fine-grained events instead of ordering among event handlers on fewer events. The resulting definition of more CTP events along the sending and receiving processing path required us to understand and interface with longer event chains when implementing new protocols. However, our experience shows that documenting and understanding the (well-defined) longer event chains was much easier than understanding somewhat shorter event chains and the ordering constraints of every possible microprotocol in the system.

V. EXPERIMENTAL RESULTS

A. Overview

While CTP cannot compete at this stage with tuned versions of TCP and UDP, the flexibility provided by the service is useful for application domains and execution environments that are not the focus of the standard protocols. In particular, CTP is useful when either a set of characteristics that falls somewhere between TCP and UDP is required, or for cases where stronger guarantees are needed than TCP provides. CTP is also appropriate when there is the opportunity to configure a protocol to match the characteristics of a specific network environment. The goal of this section is to quantify the potential overheads and benefits provided by the configurability of CTP.

In the remainder of this section, we present local area and wide area network results in a variety of situations. Local area

TABLE I
LATENCY AND BANDWIDTH COMPARISON

Protocol	Local Area		Wide Area	
	Latency (ms)	Bandwidth (Mb/sec)	Latency (ms)	Bandwidth (Mb/sec)
UDP	0.152±0.0028	90.28±0.005	413±3	61.54±1.317
CTP-Minimal	0.273±0.0052	81.76±0.014	414±0.4	56.81±0.637
CTP-LossyFIFO	0.279±0.0031	81.77±0.008	414±0.7	56.62±0.975
CTP-Video	0.366±0.0050	82.17±0.005	414±1	21.97±0.745
TCP	0.162±0.0029	89.40±0.025	412±0.9	9.38±0.189
CTP-Bulk	0.380±0.0090	68.31±0.148	415±0.8	4.73±0.288

performance results were collected between two 2-processor 2.2 GHz Pentium 3 Xeon machines running Linux kernel 2.4.18 across a quiescent 100 Mbps Ethernet; only one processor on each machine was used by the test program. The C implementation of Cactus 2.2 was used for composing microprotocols into a composite CTP protocol running at user level on top of Linux UDP sockets. Note that this imposes additional overhead on CTP compared to TCP and UDP. Wide area performance results were collected between Linux machines at the University of New Mexico (UNM) and the Georgia Institute of Technology (Georgia Tech).

Section V-B uses these platforms to quantify the cost of configurability in CTP by comparing latency and bandwidth numbers in different CTP configurations over both local and wide area networks. Section V-C then illustrates the potential benefits of CTP by customizing protocol configurations to application-specific and hardware-specific needs.

B. Configurability Overhead

The first set of experiments measures the bandwidth and ping-pong latency of UDP, TCP, and various configurations of CTP. Four different CTP configurations are included:

- CTP-Minimal: a minimal CTP configuration containing only the driver and fragmentation/reassembly microprotocols.
- CTP-LossyFIFO: the minimal CTP configuration augmented with per-message sequence numbers and unreliable in-order message delivery microprotocols.
- CTP-Video: a CTP configuration for video transmission that uses SCP-style congestion control, positive and negative acknowledgments, round-trip-time estimation, and in-order unreliable message delivery.
- CTP-Bulk: a TCP-Tahoe-like CTP configuration including reliable, in-order message delivery using retransmissions, duplicate acknowledgments, and TCP-style congestion control.

Note that the first three of these configurations are all unreliable configurations; only CTP-Bulk guarantees reliable transmission of all data.

In the latency tests, two machines ping-pong minimal-sized application packets 10 times to measure the average round-trip latency for one round trip. In the bandwidth tests, a sending application transmits 1000 1250-byte messages to a receiver, which replies with a user-level acknowledgment once all the data has been received. We measure the interval at the sender between the transmission of the first packet and receipt of the acknowledgment and use this to compute the end-to-end data

transmission rate. To enable direct comparison of protocol processing costs, the PUSH flag is set on every message handed to TCP, causing it to preserve message boundaries and send the same number of data segments as the other protocols; we confirmed experimentally that the same message boundaries were used in TCP.

Table I shows the averages and standard deviations of 10 runs of the bandwidth and latency tests on both local and wide area networks, with the top part of the table comparing unreliable protocol configurations and the bottom part comparing reliable protocols. All measurements were made on the receiver after several initial packet exchanges to allow the congestion control window to open fully.

These results indicate a latency overhead of approximately 100 microseconds per round trip over UDP in the simple local-area test and execution environment, with approximately the same service guarantees. Similarly, bandwidth is competitive with UDP, although slightly less because this version of CTP is layered on top of UDP and because of protocol overhead such as the longer CTP headers required to support the sophisticated semantics of more complex configurations (68 byte CTP headers as opposed to 8 byte UDP headers). CTP header overhead is currently unoptimized, however, and can be reduced by specializing headers to particular configurations instead of having a single generic header that encompasses all current possible CTP configurations. Additionally, running CTP directly on top of IP would lower its latency costs significantly.

As microprotocols implementing more complex semantics are added to CTP configurations in the first part of the table, latency gradually increases and bandwidth slightly decreases. Adding relatively simple microprotocols such as Lossy FIFO and Sequenced Messages to the CTP configurations (the CTP-LossyFIFO configuration) adds negligible overhead; more complex microprotocols that implement, for example, congestion control, introduce correspondingly more overhead.

In the wide area unreliable results, latencies are dominated by wide area network costs, which obscure event overhead costs. Bandwidth numbers vary as expected, with the UDP, CTP-Minimal, and CTP-LossyFIFO configurations providing the best bandwidths given their lack of congestion control. CTP-Video provides less bandwidth because of congestion control actions, but more bandwidth than the TCP and CTP-Bulk configurations. Again, this is expected since the SCP-based congestion control policy used by the multimedia-oriented CTP configurations is more aggressive than TCP-derived policies and known to not be TCP-fair.

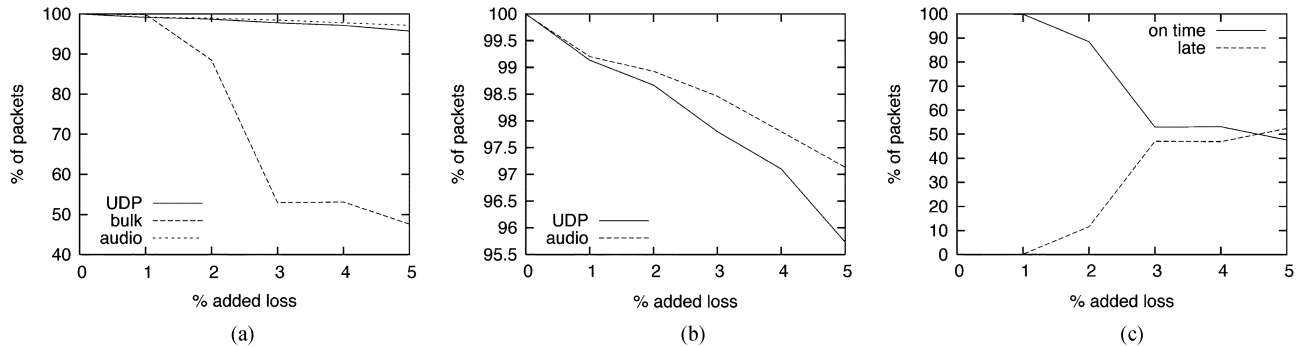


Fig. 4. Real-time streaming media performance of UDP, CTP-Bulk, and CTP-Audio. (a) UDP/CTP-Bulk/CTP-Audio Wide-Area. (b) UDP/CTP-Audio Wide-Area. (c) CTP-Bulk Breakdown.

Comparing the reliable protocols, the latency overhead of CTP-Bulk compared to TCP is somewhat higher, on the order of 200 microseconds. This is caused by the increased event processing in CTP for the complex configuration required for full reliability. We expect the latency performance of all CTP configurations to improve as the event mechanisms in the Cactus runtime are optimized, although it is probably unrealistic to expect CTP to beat TCP and UDP for this type of use.

The bandwidth differences between CTP-Bulk and TCP are caused by minor differences in delayed acknowledgment handling and packetization in the two protocol implementations. Specifically:

- CTP-Bulk currently has an MTU of 1250 bytes as opposed to the 1400+ bytes that TCP uses, has larger headers, and runs on top of UDP.
- TCP (as a stream protocol) maintains the sender window sizes in bytes, while CTP-bulk maintains a window size in packets, since it is a message-oriented protocol.

These two differences prevent CTP-Bulk from utilizing the bandwidth of a lower bandwidth wide area connection as effectively as TCP does. Note, however, that CTP's modular structure makes such differences easy to change when appropriate.

C. Benefits of Custom Configurations

CTP can be tuned to provide optimized behavior for given applications or hardware environments similar to hand-built custom protocols without the engineering overhead of developing such protocols from scratch. In this section, we demonstrate the performance benefits that customizing CTP configurations to application- and hardware-specific needs can provide.

1) *Application-Specific Customization*: To study the potential application-level benefits of protocol customization, we ran CTP as the underlying transport protocol for a custom Cactus multimedia-transmission and playback application. This application sends compressed audio or video to a remote receiver, which then plays back the received data in real-time from a playback buffer with fixed time capacity. This application supports both uncompressed and compressed (H.263/Ogg Vorbis) audio and video streams.

We studied the impact that custom CTP configurations have on an audio transmission configuration of this application using UDP, CTP-Bulk, and a new configuration CTP-Audio for audio transmission that is configured identically to CTP-Video except

for the addition of a block-erasure forward error correction microprotocol. CTP-Bulk acts as a proxy for TCP performance in this experiment, since we did not have the kernel-level access that would be needed to vary the loss experienced by the TCP protocol on the wide-area test machines. Audio packets were sent at 128 kbps on both low-latency (local) and high-latency (wide-area) networks, and with different amounts of additional packet loss at the ingress network device to examine how different protocol configurations and network conditions affected application performance. The application was set to use a fixed 3000 ms playout buffer, and CTP-Audio was set to use $N = 5$ and $K = 4$ to be able to recover from one dropped data packet out of every five packets. Each test consisted of 1500 packet transmissions, and was conducted 10 times on each protocol/network configuration.

Fig. 4(a) shows the performance of all three protocols on this application in terms of the percentage of packets delivered within the application playout window on a wide-area network between UNM and Georgia Tech. CTP-Bulk is unable to deliver packets on time in the face of significant packet loss, while UDP and CTP-Audio continue to provide reasonable service to the application. Fig. 4(b) shows only UDP and CTP-Audio performance over wide-area networks, and demonstrates that CTP-Audio is able to deliver packets on time more robustly than UDP in the face of packet loss. Local-area comparisons between UDP and CTP-Audio behave essentially the same.

Fig. 4(c) provides a more detailed breakdown of the performance of the CTP-Bulk protocol in the wide-area case. Since CTP-Bulk delivers all packets in order, as packet loss increases, packets are delivered increasingly late due to the TCP-like retransmission-based reliability scheme. UDP and CTP-Audio, on the other hand, deliver packets in a timely fashion. All of the late packets shown in 4(b) and (c) are due to packet loss, though CTP-Audio delivers more packets on time in the face of packet loss thanks to the forward-error correction service it provides to the application.

Of course, existing protocols, for example RTP [4] and SCTP [7], can provide application benefits similar to those shown above. However, each of these protocols had to be constructed from scratch, and are not easy to modify to support other, different application needs. CTP, however, allows the application authors to customize protocol behavior using a single integrated package that already supports a wide range of application-desirable semantics.

TABLE II
ROUND-TRIP LATENCY ON LOW-LATENCY UNRELIABLE NETWORKS

Loss rate (%)	Round-Trip Latency (ms)	
	CTP-Bulk	TCP
0	0.34±0.002	0.17±0.0008
1	0.62±0.07	2.1±8
2	0.66±0.07	4.4±8
3	1.9±0.5	8.0±13
4	2.1±0.4	13±18
5	2.4±0.7	16±25

2) *Hardware-Specific Customization*: In the previous case, CTP was able to be easily reconfigured to provide superior performance to applications compared to TCP and UDP because the service requirements of the application were different than those provided by TCP and UDP. However, CTP configurations can provide superior performance compared to TCP even in cases where TCP exactly matches application service requirements, particularly when the underlying network hardware violates fundamental assumptions that TCP makes.

For example, modern versions of TCP derived from the BSD code retransmit segments after receiving three duplicate ACKs or upon expiration of a retransmission timeout. However, the TCP retransmission timer is typically very coarse, on the order of 500 ms. Local wireless networks, connections across campus networks or even wide area-networks frequently yield round-trip times on the order of tens of milliseconds or faster, so faster retransmission timers can be beneficial under certain circumstances. This is particularly true on, for example, 802.11b wireless networks, which can have low latencies and high drop rates.

We have measured the performance of CTP using the CTP-Bulk configuration described above. This configuration includes the Positive ACK and Retransmit microprotocols described in Sections III-B and III-C, which use fine-grained retransmission timing. Table II lists the average round-trip latency of this CTP configuration compared to TCP. These times were measured using 10 tests of 100 back-to-back round-trips using zero-length application packets. This test was performed on the same platform as described above; network packet losses were simulated by randomly dropping varying proportions of packets on each receiving machine.

Although TCP has better latency in the lossless case, CTP was able to provide faster delivery on average when losses occurred by retransmitting more quickly. CTP can provide similar advantages in other environments where TCP is known to perform sub-optimally, such as high bandwidth-delay product links and long-distance wireless networks, or networks where losses may not be the result of congestion but may instead indicate, for example, radio interference. Moreover, CTP allows the user to configure all of these from a single integrated package, rather than forcing the construction of new specialized protocols from scratch.

D. Performance Optimizations

The performance of a composite protocol built using Cactus such as CTP can be optimized in any number of ways. These optimizations can be classified based on whether they require changes in the Cactus runtime system or microprotocols, and

the extent of these changes. The least intrusive optimizations customize the protocol's behavior using features in the Cactus runtime specifically provided for such customization. For example, message handling operations can be customized to construct message headers in whatever format is most efficient for the particular protocol by customizing the message pack and unpack routines as mentioned in Section II-B.

Another type of optimization modifies the Cactus runtime system, but does not require changes in the microprotocol code. For example, to eliminate the table lookups required to invoke customizable operations, the message handling operations can be added as static functions to the runtime system. Similarly, event dispatch and handling performance could be dramatically improved using techniques already demonstrated elsewhere [19], [20].

Finally, some optimizations require that the chosen microprotocols be modified in some way, either by hand or through automatic compile time or run time optimization. For example, the indirection required to raise an event can be optimized by replacing the raise operation with direct calls to the appropriate event handlers or even by inlining the handlers.

In the experiments above, the only optimization used was the first one described above, where the Cactus message handling operations are customized. This optimization resulted in a minor decrease in the latency and increase in the bandwidth in the CTP-Minimal and CTP-Bulk configurations over unoptimized CTP. This only took a few hours of programming and did not require any changes to the Cactus framework or CTP microprotocols. Other work has shown that more aggressive CTP optimizations can substantially improve CTP bandwidth performance [21].

VI. RELATED WORK

Other projects have explored composite protocol frameworks, generally in the context of specialized environments. Specifically, XTP [22] and TP++ [23] have been used to support flexible data transport in high-speed networks, and Minden's composite protocol system [24] supports transport protocol composition for active network systems. XTP, for example, can be configured to support different amounts of reliability and different connection establishment mechanisms. In contrast, CTP is designed to allow general configurability, enabling its use in a wide range of general purpose and specialized applications. Unlike these systems, CTP also allows the wire message format to be customized, potentially enabling backwards compatibility with protocols such as TCP and UDP.

A number of different configuration frameworks have been used to construct modular transport services [9], [25]–[28]. Most of these frameworks use a hierarchical composition model where the communication subsystem is constructed as a directed graph of modules, with interactions limited to message exchange between adjacent modules. In contrast, Cactus does not force a linear order between modules when the modules are logically on the same level or independent, and allows arbitrarily rich interactions between modules.

Protocol heaps [29] propose a non-hierarchical *role-based* approach to constructing network services and suggest that such an approach could be used either in a single layer of a protocol

stack or to replace the stack altogether. While the authors note that roles in their proposed system are similar to microprotocols in Cactus' predecessor Coyote [30], Cactus and Coyote focus only on providing support for non-hierarchical composition with a single layer of the stack. To our knowledge, protocol heaps have not yet been used to implement substantial protocols such as we have done with CTP. However, because of the similarities between the two systems, our experiences designing and implementing a flexible, non-hierarchical transport service and our measurements of the costs of such flexibility should be directly applicable to protocol heaps.

Another non-hierarchical approach is used in Adaptive [31], where each protocol or service consists of a "backplane" with slots for different protocol functions such as flow control and reliability. Unlike CTP, this backplane restricts the composition to a fixed set of functions and also constrains the interactions between different protocol functions.

Other more specialized transport protocols have been proposed since the introduction of TCP and UDP. Examples include RDP [32], which provides a message-based transport service with reliability and optional FIFO ordering guarantees, and VMTP [33], which provides transactional (RPC style) communication with customizable reliability and some support for real time and multicast data. More recent proposals include RTP [4] for transmission of real-time data such as audio or video over multicast network services, SCTP [7] for improved reliability using techniques such as multihoming, and partially-reliable transport protocols [34] for use in multimedia services. Extensions to TCP have also been developed to improve its performance and applicability for specific application or execution domains. Examples of such extensions include selective acknowledgments [35] and support for transaction-oriented services [36]. As already noted, the goal of CTP is not to be yet another transport protocol or yet another TCP extension, but a prototype of a completely customizable transport protocol that can be configured to serve any application domain in any execution environment.

VII. CONCLUSION

The ability to customize transport protocols can provide important flexibility when it comes to supporting new applications and new network technologies. Here, we have described an approach to building such services based on Cactus, as well as a concrete realization of the approach in the form of CTP. In this family of transport protocols, various attributes are implemented as separate microprotocols, which are then combined in different ways to provide customized semantics. Initial experimental results indicate that, while the performance is somewhat slower than TCP and UDP for similar configurations, the ability to target the guarantees more precisely can in fact result in better performance. While it will always be possible to construct more efficient specialized solutions, CTP allows easy component-based construction of custom transport protocols with minimal effort.

Future work will focus on using CTP as an experimentation and prototyping platform to implement and measure different transport-related algorithms. This will require extending CTP with, for example, MPI-style matching instead of port-based demultiplexing. We also plan to extend CTP to support customiz-

able multicast and group communication, as well as to explore further performance optimizations.

ACKNOWLEDGMENT

The authors would like to thank the referees, who provided valuable comments that substantially improved the paper.

REFERENCES

- [1] S. Kent and R. Atkinson, "Security architecture for the Internet Protocol," RFC (Standards Track) 2401, 1998.
- [2] A. Freier, P. Karlton, and P. Kocher, "The SSL protocol, version 3.0," Netscape Communications, Internet-Draft, 1996.
- [3] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: a new resource ReSerVation Protocol," *IEEE Network*, vol. 7, no. 5, pp. 8–18, Sep. 1993.
- [4] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," RFC 1889, 1996.
- [5] R. Wu and A. Chien, "GTP: Group transport protocol for lambda-grids," in *Proc. IEEE Int. Symp. Cluster Computing and the Grid (CC-Grid 2004)*, Chicago, IL, Apr. 2004, pp. 228–238.
- [6] E. Weigle and A. Chien, "The Composite Endpoint Protocol (CEP): scalable endpoints for terabit flows," in *Proc. IEEE Int. Symp. Cluster Computing and the Grid (CCGrid 2005)*, Cardiff, U.K., May 2005, pp. 1126–1134.
- [7] R. R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. J. Schwarzbauer, T. Taylor, I. Rytina, M. Kalla, L. Zhang, and V. Paxson, "Stream Control Transmission Protocol," Internet Draft, 2000.
- [8] M. Hiltunen, R. Schlichting, X. Han, M. Cardozo, and R. Das, "Real-time dependable channels: customizing QoS attributes for distributed systems," *IEEE Trans. Parallel Distrib. Syst.*, vol. 10, no. 6, pp. 600–612, Jun. 1999.
- [9] N. Hutchinson and L. Peterson, "The x-kernel: An architecture for implementing network protocols," *IEEE Trans. Software Eng.*, vol. 17, no. 1, pp. 64–76, Jan. 1991.
- [10] M. Hiltunen, "Configuration management for highly-customizable software," *IEE Proc.: Software*, vol. 145, no. 5, pp. 180–188, 1998.
- [11] W.-K. Chen, M. Hiltunen, and R. Schlichting, "Constructing adaptive software in distributed systems," in *Proc. 21st Int. Conf. Distributed Computing Systems*, Mesa, AZ, 2001, pp. 635–643.
- [12] P. Bridges, "Composing and coordinating adaptation in Cholla," Ph.D. dissertation, Univ. Arizona, Tucson, AZ, 2002.
- [13] J. Appavoo, K. Hui, C. Soules, R. Wisniewski, D. D. Silva, O. Krieger, M. Auslander, D. Edelsohn, B. Gamsa, G. Ganger, P. McKenney, M. Ostrowski, B. Rosenburg, M. Stumm, and J. Xenidis, "Enabling autonomous system software with hot-swapping," *IBM Syst. J.*, vol. 42, no. 1, pp. 60–76, 2003.
- [14] L. Rizzo, "Effective erasure codes for reliable computer communication protocols," *Comput. Commun. Rev.*, vol. 27, no. 2, pp. 24–36, 1997.
- [15] V. Jacobson, "Congestion avoidance and control," in *Proc. ACM SIGCOMM'88*, Stanford, CA, Aug. 1988, pp. 314–332.
- [16] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based congestion control for unicast applications," in *Proc. ACM SIGCOMM 2000*, Stockholm, Sweden, Aug. 2000, pp. 43–56.
- [17] S. Cen, C. Pu, and J. Walpole, "Flow and congestion control for Internet streaming applications," presented at the ACM/SPIE Multimedia Computing and Networking Conf. (MMCN98), San Jose, CA, Jan. 1998.
- [18] G. T. Wong, M. A. Hiltunen, and R. D. Schlichting, "A configurable and extensible transport protocol," in *Proc. IEEE INFOCOM 2001*, Anchorage, AK, Apr. 2001, pp. 319–328.
- [19] M. Rajagopalan, S. Debray, M. Hiltunen, and R. Schlichting, "Profile-directed optimization of event-based programs," in *Proc. ACM SIGPLAN 2002 Conf. Programming Language Design and Implementation*, 2002, pp. 106–116.
- [20] C. Chambers, S. Eggers, J. Auslander, M. Philipose, M. Mock, and P. Pardyak, "Automatic dynamic compilation support for event dispatching in extensible systems," presented at the 1st Workshop on Compiler Support for Systems Software (WCSS'96), Tucson, AZ, Feb. 1996.
- [21] R. Wu, A. Chien, M. Hiltunen, R. Schlichting, and S. Sen, "A high performance configurable transport protocol for Grid computing," in *Proc. IEEE Int. Symp. Cluster Computing and the Grid (CCGrid 2005)*, Cardiff, U.K., May 2005, pp. 1117–1125.
- [22] "Xpress transport protocol specification, revision 4.0," XTP Forum, Santa Barbara, CA, Mar. 1995.
- [23] D. Feldmeier, "An Overview of the TP++ Transport Protocol Project," in *High Performance Networks, Frontiers and Experience*. Boston, MA: Kluwer Academic, 1994, pp. 157–176.

- [24] G. Minden, E. Komp, S. Ganje, M. Kannan, S. Subramaniam, S. Tan, S. Vallabhaneni, and J. Evans, "Composite protocols for innovative active services," in *Proc. 2002 DARPA Active Networks Conf. and Expo.*, May 2002, pp. 157–165.
- [25] D. Ritchie, "A stream input-output system," *AT&T Bell Labs Tech. J.*, vol. 63, no. 8, pp. 311–324, 1984.
- [26] R. van Renesse, K. Birman, M. Hayden, A. Vaysburd, and D. Karr, "Building adaptive systems using ensemble," *Software Practice and Experience*, vol. 28, no. 9, pp. 963–979, 1998.
- [27] W. Allcock, J. Bresnahan, R. Kettimuthu, and J. M. Link, "The Globus eXtensible Input/Output System (XIO): A protocol independent IO system for the Grid," in *Proc. 19th IEEE Int. Parallel and Distributed Processing Symp. (IPDPS 2005)*, Denver, CO, Apr. 2005, 8 pp..
- [28] G. Parr and K. Curran, "A paradigm shift in the distribution of multimedia," *Commun. ACM*, vol. 43, no. 6, pp. 103–109, 2000.
- [29] R. Braden, T. Faber, and M. Handley, "From protocol stack to protocol heap—Role-based architecture," presented at the 1st Workshop on Hot Topics in Networks (HotNets-I), Princeton, NJ, Oct. 2002.
- [30] N. Bhatti, M. Hiltunen, R. Schlichting, and W. Chiu, "Coyote: A system for constructing fine-grain configurable communication services," *ACM Trans. Comput. Syst.*, vol. 16, no. 4, pp. 321–366, 1998.
- [31] D. Schmidt, D. Box, and T. Suda, "ADAPTIVE: A dynamically assembled protocol transformation, integration, and evaluation environment," *Concurrency: Practice and Experience*, vol. 5, no. 4, pp. 269–286, 1993.
- [32] D. Velten, R. Hinden, and J. Sax, "Reliable Data Protocol," RFC 908, 1984.
- [33] D. Cheriton, "VMTP: Versatile Message Transaction Protocol," RFC 1045, 1988.
- [34] R. Marasli, P. Amer, and P. Conrad, "Partially reliable transport service," in *Proc. 2nd IEEE Symp. Computers and Communications*, 1997, pp. 648–656.
- [35] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "TCP selective acknowledgment options," RFC 2081, 1996.
- [36] R. Braden, "T/TCP—TCP extensions for transactions," RFC 1644, 1994.



Patrick G. Bridges (M'03) received the B.S. degree from Mississippi State University, State College, MS, and the Ph.D. degree from the University of Arizona, Tucson, AZ.

He is an Assistant Professor with the Computer Science Department at the University of New Mexico, Albuquerque, NM. His research interests include high-performance computing systems, configurable and adaptable system software, and system-wide measurement and monitoring techniques.

Dr. Bridges is a member of the ACM and the IEEE Computer Society.



Gary T. Wong received the B.S. degree from the University of Auckland, New Zealand.

He is a researcher with the Computer Science Department at Boston University, Boston, MA, where he collaborates with many others on a video sensor network project. His past work covers a range of topics from compilers to distributed systems, but he is particularly interested in exploring novel structures for operating systems.



Matti A. Hiltunen (M'07) received the M.S. degree from the University of Helsinki, Finland, and the Ph.D. degree from the University of Arizona, Tucson, AZ.

He is a researcher in the Dependable Distributed Computing and Communication department at AT&T Labs—Research, Florham Park, NJ. His research interests include dependable distributed systems and networks, grid computing, and pervasive computing.

Dr. Hiltunen is a member of the ACM, the IEEE Computer Society, and IFIP Working Group 10.4.



Richard D. Schlichting (F'02) received the B.A. degree from the College of William and Mary, Williamsburg, VA, and the M.S. and Ph.D. degrees from Cornell University, Ithaca, NY.

He is Executive Director of Software Systems Research at AT&T Labs—Research, Florham Park, NJ. His research interests include distributed systems, highly dependable computing, and networks.

Dr. Schlichting is an ACM Fellow and an IEEE Fellow, and is the current chair of IFIP Working Group 10.4.



Matthew J. Barrick (S'07) received the B.S. degree from the Indiana University of Pennsylvania, Indiana, PA. He is pursuing the Ph.D. degree in the Computer Science Department at the University of New Mexico, Albuquerque, NM.

He currently works in the UNM Scalable Systems Lab on issues ranging in a variety of areas, including high-performance large-scale file systems and scalable multiprocessor operating systems like K42.

Mr. Barrick is a member of the ACM.