A Taxonomy and Survey of SCTP Research

ŁUKASZ BUDZISZ, Technische Universität Berlin
JOHAN GARCIA and ANNA BRUNSTROM, Karlstad University
RAMON FERRÚS, Universitat Politècnica de Catalunya

The Stream Control Transmission Protocol (SCTP) is a relatively recent general-purpose transport layer protocol for IP networks that has been introduced as a complement to the well-established TCP and UDP transport protocols. Although initially conceived for the transport of PSTN signaling messages over IP networks, the introduction of key features in SCTP, such as multihoming and multistreaming, has spurred considerable research interest surrounding SCTP and its applicability to different networking scenarios. This article aims to provide a detailed survey of one of these new features—multihoming—which, as it is shown, is the subject of evaluation in more than half of all published SCTP-related articles. To this end, the article first summarizes and organizes SCTP-related research conducted so far by developing a four-dimensional taxonomy reflecting the (1) protocol feature examined, (2) application area, (3) network environment, and (4) study approach. Over 430 SCTP-related publications have been analyzed and classified according to the proposed taxonomy. As a result, a clear perspective on this research area in the decade since the first protocol standardization in 2000 is given, covering both current and future research trends. On continuation, a detailed survey of the SCTP multihoming feature is provided, examining possible applications of multihoming, such as robustness, handover support, and loadsharing.

Categories and Subject Descriptors: C.2.2 [**Computer-Communication Networks**]: Network Protocols; C.2.6 [**Computer-Communication Networks**]: Internetworking

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: SCTP, taxonomy, multihoming, robustness, handoff management, loadsharing

ACM Reference Format:
Budzisz, Ł., Garcia, J., Brunstrom, A., and Ferrús, R. 2012. A taxonomy and survey of SCTP research. ACM Comput. Surv. 44, 4, Article 18 (August 2012), 36 pages.
DOI = 10.1145/2333112.2333113 http://doi.acm.org/10.1145/2333112.2333113

1. INTRODUCTION

At the inception of IP-based networks in the early 1970s, telephony networks and computer communication networks were viewed as quite different entities. The operational requirements, technical design choices, and standardization procedures were some of the differentiating factors at the time. In the last decade and a half, there has been

This work was initiated with support from the EU network of excellence NEWCOM (ref. FP6-IST-507325). This work was also partly supported by the Spanish Research Council and FEDER funds under ODYSSEY (TEC2006-09109) and ARCO grants (TEC2010-15198).
Authors' addresses: Ł. Budzisz, Telecommunication Networks Group (TKN), Technische Universität Berlin (TU Berlin), Einsteinufer 25, 10587 Berlin, Germany; email: lukasz.budzisz@tu-berlin.de; J. Garcia and A. Brunstrom, Department of Computer Science, Karlstad University, Universitetsgatan 2, 651 88 Karlstad, Sweden; email: {Johan.Garcia,Anna.Brunstrom}@kau.se; R. Ferrús, Mobile Communications Research Group (GRCM), Department of Signal Theory and Communication (TSC), Universitat Politècnica de Catalunya (UPC), C/. Jordi Girona 1–3, 08034 Barcelona, Spain; email: ferrus@tsc.upc.edu.

DOI 10.1145/2333112.2333113 http://doi.acm.org/10.1145/2333112.2333113

an increasing convergence between the once quite separate worlds of telephony and computer networking. This convergence has had many effects, some quite drastic while others have been more behind the scenes. This article relates to one noteworthy effect of this convergence, that is, the need to transport telephony signaling over IP networks. A major enabler for this is the Stream Control Transmission Protocol (SCTP), a protocol originally developed to carry the Signaling System No. 7 (SS7) signaling traffic found in the traditional public switched telephone network (PSTN) over IP networks.

The first SCTP specification was published in October 2000 by the Internet Engineering Task Force (IETF) Signaling Transport (SIGTRAN) working group[1] in the now obsolete RFC 2960 [Stewart et al. 2000]. Since then, the original protocol specification has been slightly modified (checksum change, RFC 3309 [Stone et al. 2002]) and updated with suggested implementer's fixes (RFC 4460 [Stewart et al. 2006]). Both updates are included in the current protocol specification, RFC 4960 [Stewart 2007] that was released in September 2007. Already at an early stage of the specification work, it was envisaged that the capabilities of SCTP would make it suitable as a general transport protocol.[2] Indeed, the following years saw a growing range of possible applications of SCTP, with many works discussing both signaling and more general purposes. A considerable amount of research has been generated examining both the additional functionality per se as well as the challenges and possibilities created by a general transport protocol setting for SCTP.

Consequently, this substantial amount of published research on SCTP requires some categorization to provide a manageable overview of current state of the art. To this end, we have developed a taxonomy which is used to classify the research along four orthogonal dimensions: protocol feature examined, application area, network environment, and study approach. The presented taxonomy allows us to analyze SCTP research trends over time and evaluate the impact of the protocol in various areas. To illustrate the usefulness of the provided classification apart from a general research analysis, we chose one of the protocol features (multihoming) to conduct a detailed survey of the collected research. Moreover, the research analysis provided in this article has also led to the development of an extensive database, available online [Garcia and Budzisz 2010], where SCTP-related references can be explored according to the dimensions and categories of the presented taxonomy.

The bulk of the available SCTP-related research used to provide an organized classification presented in this report has been collected from seven of the most common bibliographic databases.[3] The time frame for the collected research encompasses ten years since the publication of RFC 2960 (i.e., 2000 is the initial year when SCTP research took off, whereas the end of 2009 has been set up as the cut-off date in order to guarantee a trustworthy metric for this article). Figure 1 shows the annual distribution of published SCTP–related articles in the discussed time frame. Out of 434 total articles collected, merely 15% were published in the first four years since the protocol specification was announced. Then, a stable increasing trend can be observed, with the

[1] `http://tools.ietf.org/wg/sigtran/`.

[2] The fourth draft of the Multi-network Datagram Transmission Protocol (MDTP), an ancestor of SCTP dating back to April 1999, removed the limitation to signaling transport.

[3] The following databases were taken into account: (1) the IEEE Xplore database (`http://ieeexplore.ieee.org/xplore/dynhome.jsp`); (2) the ACM Digital Library (`http://portal.acm.org/`); (3) the BibFinder database (`http://kilimanjaro.eas.asu.edu/`); (4) the Engineering Village database (`http://www.engineeringvillage2.com/`); CiteSeer.IST, CiteSeer, IST, Scientific Literature Digital Library (`http://citeseer.ist.psu.edu/`); Google™Scholar (`http://scholar.google.com/`); and (7) the ISI Web of Knowledge (`http://www.isiwebofknowledge.com/`).

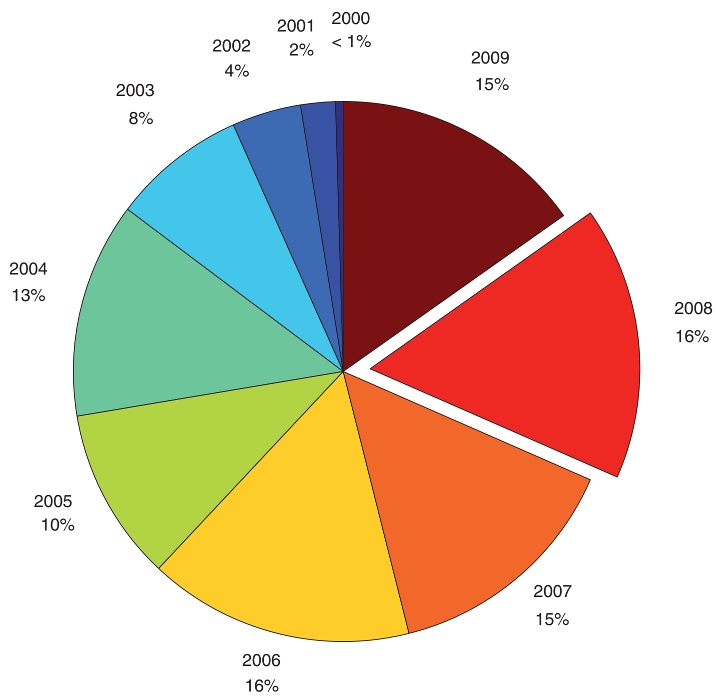

Fig. 1. Annual distribution of all published articles.

largest number of articles (71) published in 2008; however, the differences between any of the last six years are relatively small.

Despite a considerable amount of research, it is worth noting that SCTP still lacks a killer application that could motivate its widespread adoption into the well-established IP networks protocol stack. Hence, after more than ten years since the appearance of the first protocol specification, SCTP is still not part of the vendor-supplied TCP/IP stack for widespread OSes, such as MS Windows™or Mac OS X™, and the endpoint support for communication services and applications is quite weak so far. Nevertheless, an important milestone towards a broader adoption of SCTP was the decision within the mobile communications industry to select SCTP as a transport protocol for the Long Term Evolution (LTE) networks to support signaling message exchange between network nodes. SCTP is also the key transport component in current SIGTRAN suites used for transporting SS7 signaling information over packet-based networks. Hence, SCTP is used in progressively adopted Voice over IP (VoIP) architectures and thus becomes part of related signaling gateways, media gateway controllers, and IP-based service control points that are used to develop convergent voice and data solutions.

The remainder of the article is structured as follows. In Section 2, a brief general description of SCTP is provided. Next, Section 3 explains the methodology for creating the SCTP research taxonomy and describes the taxonomy dimensions and categories. In Section 4, the taxonomy is used on the available SCTP-related research material to analyze and provide trends on the distribution across the proposed dimensions. Section 5 presents a detailed survey of the key research findings related to the SCTP multihoming feature and its main applications: robustness, handover support, and loadsharing. Finally, conclusions are drawn in Section 6.

Table I. Comparison of Transport-Layer Protocols

Feature	SCTP	TCP	UDP
Connection-oriented	Yes	Yes	No
Half-closed connections	No	Yes	N/A
Protection against blind DoS attacks	Yes	No	N/A
Dynamic address manipulation	Optional[a]	No	N/A
Reliable data transfer	Yes	Yes	No
Partially reliable data transfer	Optional[b]	No	No
Preservation of application message boundaries	Yes	No	Yes
Application PDU fragmentation/bundling	Yes	Yes	No
Ordered data delivery	Yes[c]	Yes	No
Unordered data delivery	Yes	No	Yes
Full-duplex data transmission	Yes	Yes	Yes
Flow and congestion control	Yes	Yes	No
Selective acknowledgments	Yes	Optional	No
Path max. transmission unit discovery	Yes	Yes	No
Explicit congestion notification support	Yes	Yes	No
Multistreaming	Yes	No	No
Multihoming	Yes	No	No

[a] Covered with DAR extension, see RFC 5061 [Stewart et al. 2007].
[b] Covered with PR-SCTP extension, see RFC 3758 [Stewart et al. 2004].
[c] The data within a stream is delivered in order.

2. SCTP BACKGROUND

In its most basic use case, SCTP provides a transport service quite similar to that of the Transmission Control Protocol (TCP), that is, a reliable full-duplex connection with flow and congestion control. Many aspects of SCTP, such as the flow and congestion control algorithms, are based on the corresponding TCP functionality and are thus quite similar to its widely used ancestor, as further discussed in Section 2.1. At the same time, SCTP aims to solve some important and well-known problems of TCP by bringing new key features into the transport layer. These features are described in more details in Section 2.2, introducing multistreaming[4] and multihoming. As shown in the later part of this article, they are the two features that have made SCTP a subject of considerable research.

Table I, first provided in Stewart and Amer [2007], is extended here to highlight the main characteristics of SCTP along with those of TCP and the User Datagram Protocol (UDP). Good tutorials covering the basics of SCTP can be found in many works [Stewart and Metz 2001; Barile 2004; Fu and Atiquzzaman 2004]. Such a detailed explanation is out of the scope of the current article whose focus is particularly aimed at analyzing SCTP-related research. In any case, to facilitate the readability of the next sections and achieve a level of self containment, some SCTP basics and new features are briefly described next.

[4] The feature from which the SCTP name is actually derived.

2.1. Protocol Basics

The SCTP protocol data unit (SCTP-PDU), simply called the *SCTP packet*, consists of an SCTP common header and one or more *chunks*. A chunk carries chunk-specific information, being either control or user data. Several chunks, except those of types INIT, INIT ACK, and SHUTDOWN COMPLETE which are used in the establishment and termination of a connection, may be bundled into one SCTP packet up to the limit given by the maximum transmission unit (MTU) size.

SCTP is a connection-oriented protocol that works on top of the connectionless IP network. The term *association* is used to name the relationship between two SCTP endpoints, distinguishing it from a TCP connection as a more complex structure that can span over multiple IP addresses at each endpoint. An SCTP association is set up in a four-way handshake, where a sequence of INIT, INIT-ACK, COOKIE-ECHO, and COOKIE-ACK chunks are exchanged between the peers. This is one of the differences compared to TCP, which uses a three-way handshake prone to blind denial of service attacks, such as SYN flooding, that unnecessarily reserves server resources for bogus connections. To prevent this SCTP has a cookie mechanism in which a cookie is created using a secret key and a hash mechanism, adding an additional leg to the association setup before the actual resource reservation can take place. To enable the association and reserve the resources, the initiator must answer with a COOKIE-ECHO containing the same cookie as received in the INIT-ACK. At the end, the COOKIE-ACK is sent back to the initiator to acknowledge the association setup. Another difference compared to TCP can be seen at the release of an association, which is simpler in SCTP and involves only a three-way handshake (SHUTDOWN, SHUTDOWN-ACK, and SHUTDOWN-COMPLETE). As a consequence, and in contrast to TCP, SCTP does not allow half-closed connections.

SCTP's flow and congestion control algorithms are essentially the same as those in TCP. SCTP provides flow control to prevent the sender from overflowing the receiver's buffer. Identically as in TCP, the SCTP sender maintains an association variable called the *advertised receiver window size (a_rwnd)* to keep track of the space that is currently available in the receiver buffer. Initial a_rwnd credit is announced in the INIT and INIT ACK chunks and is updated in each confirmation of successful data reception (SACK).[5] According to the flow control algorithm, the sender can not send any new data if the receiver indicates that it has no space to buffer that data (i.e., a_rwnd is 0). Regardless of the a_rwnd value, the sender can have one DATA chunk in flight to provide information about rwnd changes that might have been missed due to the last SACK being lost. Relative to TCP, the fact that SCTP has multistreaming and can be configured for unordered delivery may affect how SCTP's a_rwnd calculation is done.

Congestion control prevents the SCTP sender from overloading the network. The basic idea is to drastically reduce the sending rate at the event of packet loss, which is considered as an indication of congestion. Consequently, when there is no congestion event, the SCTP sender *additively increases* its sending rate, and once the congestion is detected, that is, a loss occurs, the SCTP sender *multiplicatively decreases* its sending rate. This approach, called *AIMD behavior*, is achieved using four main algorithms that both TCP and SCTP employ: slow start, congestion avoidance, fast retransmission, and fast recovery. TCP congestion control state of the art for the late 1990s that served as a base for SCTP is described in Paxon et al. [1999]. As mentioned in the SCTP protocol specification [Stewart 2007], one of the few modifications relative to TCP is that SCTP adds one more control variable for regulating its sending rate. Apart from the congestion window (cwnd) and slow-start threshold (ssthresh) known from TCP, there

[5]SCTP uses a selective acknowledgment mechanism (SACK), which is optional for TCP.

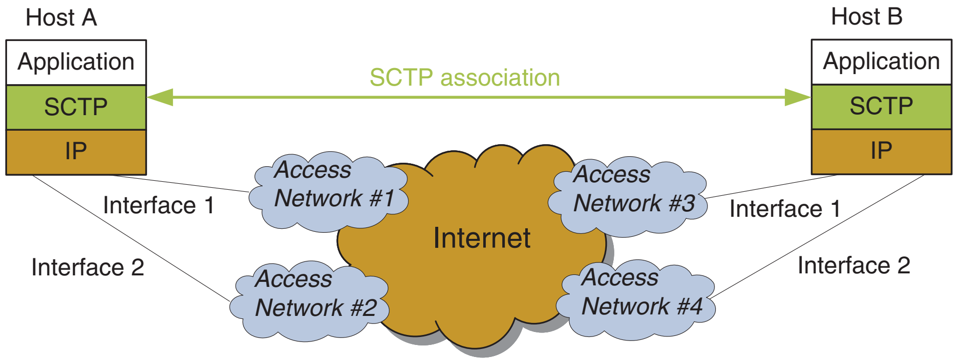

(a) Multihoming.

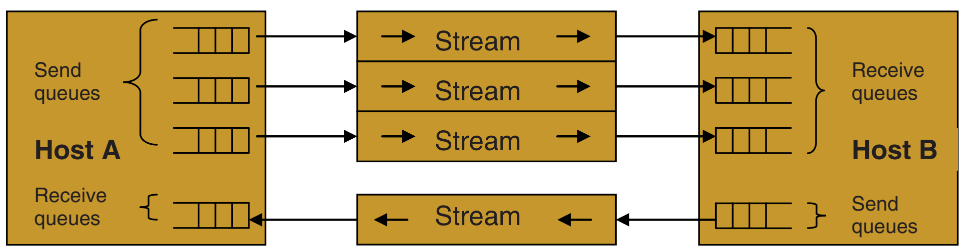

(b) Multistreaming.

Fig. 2. New SCTP features.

is also a variable (partial_bytes_acked) that keeps track of all data acknowledged (not necessarily in sequence) from the last cwnd increase. This new variable is only used in the congestion avoidance phase in order to improve the accuracy of cwnd adjustments.

2.2. New Protocol Features

One of the features that SCTP supports is *multihoming*. As illustrated in Figure 2(a), multihoming allows a single association between two SCTP endpoints to combine multiple source and destination IP addresses. These IP addresses are exchanged and verified during the association setup, and each destination address is considered a different path towards the corresponding endpoint. Also during the association setup, one path among all available is selected as the *primary* path (by default the source-destination pair where the association is established), and the others are designated as *alternate* paths. Alternate paths are often referred to in the literature as *backup* paths, especially in the robustness context of multihoming. The availability of the paths is verified by means of special control chunks called HEARTBEATs that are sent periodically to all destinations. As to the standardized behaviour in RFC 4960, data transmission is always conducted through the primary path, while backup paths are only used to handle data retransmissions. A temporary congestion or permanent failure of the primary path can cause the protocol to shift data transmission over to one of the backup paths. The change of the primary path is referred to as a failover mechanism. Multihoming was originally designed to provide increased robustness to applications requiring high availability, such as the SS7 signaling transport. Despite the evolution of SCTP towards a general transport protocol, this design principle has been kept in the current specification. Other applications of the multihoming feature, such as transport layer mobility or load balancing over multiple network paths, are not explicitly supported within the standard SCTP specification.

An important extension to SCTP in the context of its multihoming feature is the dynamic address reconfiguration (DAR) extension, defined in RFC 5061 [Stewart et al. 2007]. DAR makes it possible to dynamically add or delete IP addresses and to request a primary-path change during an active SCTP association. Although originally defined to help with IPv6 renumbering and hot-pluggable cards, the DAR extension can be easily leveraged to make SCTP a mobility enabled transport protocol. All previously mentioned applications of multihoming are surveyed in detail in Section 5, since they represent an important share of the conducted research. Finally, it is worth mentioning that multihoming typically, but not necessarily, implies multiple network interfaces in the SCTP endpoints. Indeed assignment of multiple IP addresses to the same network interface suffices to exploit SCTP multihoming.

Multistreaming, illustrated in Figure 2(b), is another new SCTP feature that allows for the establishment of associations with multiple streams. Streams are unidirectional data flows within a single association. The number of requested streams between peer SCTP endpoints is declared during the association setup, and the streams are valid during the entire association lifetime. Each stream is distinguished by the stream identifier field included in each data chunk so that chunks from different streams can be bundled inside one SCTP packet. To preserve order within a stream, the stream sequence number is used. In such a case, when a transmission error occurs on one of the streams, it does not affect data transmission on the other streams. Consequently, TCP's head-of-line (HoL) blocking problem is reduced to the affected stream only, and does not stall the entire association. Notice that when a transmission loss occurs on a TCP connection, packet delivery is suspended until the missing parts are restored, as in-sequence data delivery is a key TCP feature. This may cause additional data transmission delay to all data sent over the TCP connection.

Among the most important applications of multistreaming are priority stream scheduling [Heinz and Amer 2004], preferential treatment [Samtani et al. 2003], and reducing the latency of streaming multimedia in high-loss environments [Kim et al. 2007; Sheu and Tu 2007]. There is also the SCTP partially reliable extension (PR-SCTP), standardized in RFC 3758 [Stewart et al. 2004], which offers a non-duplicate data delivery service with tunable loss recovery. Together with the multistreaming feature, PR-SCTP can be used to provide better support to real-time applications.

The aforementioned array of new features that SCTP offers has attracted researchers from diverse fields and created many new ideas. To increase the ability to discuss this new body of research in a structured manner, a taxonomy is presented next.

3. SCTP RESEARCH TAXONOMY

The purpose of a taxonomy is, in general terms, to provide a classification based upon a given scheme in order to obtain a comprehensive understanding of the addressed topic. How the taxonomy scheme should be constructed is an open question and must be tailored to each specific instance. Ideally, the scheme should be complete so that all works fit within the taxonomy, and the classification categories should be non-overlapping with well-defined limits between the categories. The taxonomy proposed here for analyzing SCTP-related research is constructed using four orthogonal dimensions with a number of non-overlapping categories in each dimension. The dimensions represent the different aspects to be analyzed for each classified research article related to SCTP. The four dimensions used for classification are (1) protocol feature examined, (2) application area, (3) network environment, and (4) study approach.[6] Each dimension consists of a set of categories used to classify all SCTP-related articles. A given article may belong

[6]The selected dimensions are quite generic and may also be applicable in taxonomy schemes covering other networking research areas.

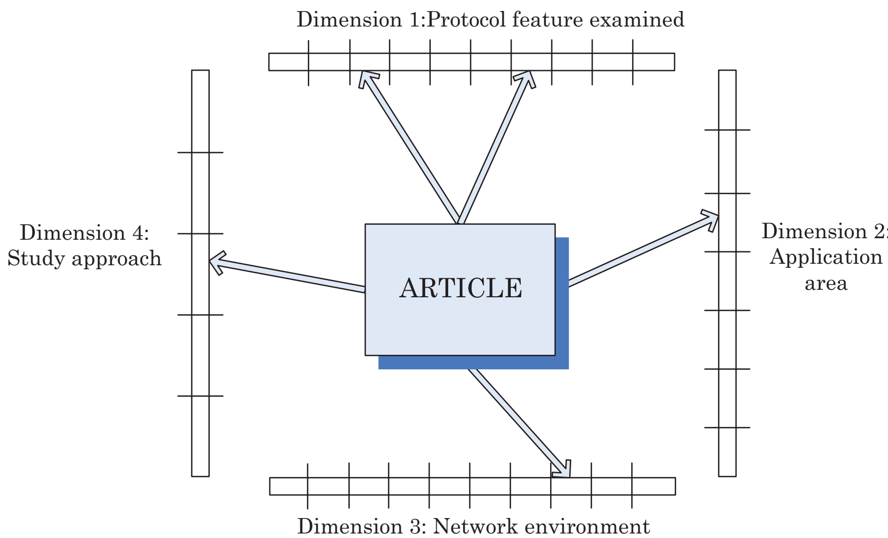

Fig. 3. Graphical illustration of proposed taxonomy.

to one or more categories per dimension. For example, an article may examine several of the SCTP protocol features or use multiple study approaches. The classification of an SCTP-related article based on the proposed taxonomy is illustrated in Figure 3.

The main design goal and, unsurprisingly, the hardest problem to tackle at the early stage of defining the classification categories was to minimize the possible overlap between them in order to reduce as much as possible the ambiguity of which category some research aspect may relate to. As classification work progressed, the initial set of proposed categories was refined to create a final categorization that has a high degree of orthogonality and is versatile enough to evaluate practically all SCTP-related research. This iterative mode of taxonomy development also resulted in a reduced complexity for the final taxonomy that will now be presented in detail.

3.1. Dimension 1: Protocol Feature Examined

The first dimension in this taxonomy classifies the research into different categories defined upon the protocol feature or functionality examined. As pointed out in Section 2, compared to TCP, SCTP provides new interesting functionality. Much of the SCTP research obviously targets this new functionality and analyzes it from different viewpoints. Besides that, SCTP has a number of features, identical or similar to TCP, that have also spurred some research. For the case of multihoming, a feature investigated by more than half of the classified articles, it was considered appropriate to distinguish three separate subcategories depending on the aim for which multimoming is used. The categories in Dimension 1 are the following.

(1) *MH-Robust*. The multihoming feature was originally designed for enhancing end-to-end robustness by using transport layer failover to an alternate path when the primary path fails.
(2) *MH-Handover*. The multihoming functionality of SCTP can also be used as a building block to provide transport layer mobility management solutions.
(3) *MH-Load Sharing*. The multihoming feature may be used to concurrently transfer data over multiple paths in a load-balancing fashion. This creates both a potential

for improved end-to-end performance and a number of complicating issues that need to be addressed.

(4) *DAR*. The extension to SCTP for dynamic address manipulation during an active SCTP association is an essential building block for handling mobility at the transport layer for either single- or multi-homed nodes.

(5) *Multistreaming*. The multistreaming capability of SCTP allows a single association to have multiple logically separate streams. This functionality is new relative to TCP, and one major advantage is that it can reduce the HoL blocking that may occur with TCP, since there is no ordering requirement between the streams.

(6) *PR-SCTP*. The SCTP partially reliable extension offers more flexibility with regards to the reliability of the transport service. The ability to provide partial reliability opens up new possibilities in how to handle reliability-versus-latency trade-offs at various layers in the protocol stack.

(7) *SH-Congestion*. Congestion control is a central issue for any transport protocol that is to be deployed on the Internet. SCTP's congestion control for singlehomed associations is to a large extent similar to TCP's but also has some differences, as mentioned in Section 2.1. There is an abundance of literature on TCP congestion control, and the impact of SCTP's congestion control nuances on wired and wireless networks is clearly a relevant topic.

(8) *Security*. Research on security mainly addresses two aspects. First, SCTP provides some new security enhancements for improving the main vulnerabilities of TCP. Second, the multihoming feature together with the DAR extension permits address manipulation and raises additional security concerns, such as hijacking attacks.

(9) *Survey*. This specific category captures all publications that provide a general protocol overview rather than an analysis of a particular protocol feature.

(10) *Other*. Some additional and diverse aspects of SCTP that are not covered by any of the preceding categories (i.e., checksum usage, SCTP sockets, and some experimental extensions of the protocol architecture) are classified here. For the sake of avoiding an excessive number of categories in this dimension, it has been decided that the definition of a new category requires at least ten papers addressing a given aspect.

3.2. Dimension 2: Application Area

The second dimension focuses on the application area adressed by the research to be classified. The application area influences both the characteristics of the traffic data to be transferred with SCTP (e.g., message size distribution, data generation process, etc.) and the required transmission performance targets. The categories in Dimension 2 are the following.

(1) *Signaling*. Since SCTP was originally designed for transporting SS7 signaling, performance analysis in this application domain is important. SCTP is not limited to SS7 signaling but can also be used to transfer other kinds of signaling traffic, such as Session Initiation Protocol (SIP).

(2) *Multimedia*. SCTP, especially together with the PR-SCTP extension that provides partial reliability, can be used to transfer multimedia data. Also, the multistreaming capabilities of SCTP map very well to multimedia traffic having multiple media streams.

(3) *Web*. Web transfer is a large application area for TCP and also may be so for SCTP. SCTP's multistreaming ability is one of the factors that may impact the transport layer performance of SCTP in this application area.

(4) *Bulk*. Applications such as the File Transfer Protocol (FTP) are commonly used when examining transport layer protocols, and these are also relevant for SCTP. Examination of singlehomed SCTP for bulk transfer, for example, provides insights into the steady-state protocol performance and allows comparison to similar TCP results. Bulk traffic is characterized by being large enough to have the transfer time decided by the steady-state behavior and by being greedy, i.e., always having data to send when there are transmission resources available.
(5) *MPI*. The Message Passing Interface (MPI) application area covers the use of SCTP in local and wide area cluster and grid environments. MPI is now commonly used in high-performance computing, and SCTP is seen as a promising option of IP-based transport support for MPI.
(6) *Other-Applications*. This category groups a few other specified (but diverse) applications, not captured by any of the above categories, such as for example data acquisition systems (DAQ). None of the applications grouped into this category has more than 5 papers.
(7) *Unspecified*. This category is used when no application area has been defined in the classified research, such as research focusing on conceptual discussions of some aspect.

3.3. Dimension 3: Network Environment

The third dimension covers the network environment that is considered in the research to be classified. This dimension is divided into ten categories, grouped into the domains *wireless* and *wired* to improve clarity. The categories in Dimension 3 are the following.

Wireless domain.

(1) *WLAN*. The wireless local area network (WLAN) environment is characterized by relatively high bandwidth and a reliable link layer that, to a large extent, shields the upper layers from physical layer problems, such as bit errors and frame losses, but can show significant end-to-end delay variation.
(2) *MANET / VANET*. The mobile ad hoc network (MANET) environment has a number of defining characteristics, such as relatively low bandwidth, large variation in end-to-end delays, and losses directly caused by congestion or temporary route unavailability. This category also includes vehicular ad hoc networks (VANET), since they are similar in nature and there are too few articles in the latter group to motivate a separate category.
(3) *Cellular*. This category, which spans from 2G mobile networks with limited data transfer support to beyond 3G high-speed packet access (e.g., 3GPP HSPA), offers low to medium bandwidths and considerable variation in end-to-end delays. In these types of networks, a reliable link layer typically ensures that there will be no wireless losses, only congestion losses.
(4) *Heterogeneous*. This category refers to the case in which coexisting heterogeneous wireless networks are considered in the research, which is typical for vertical handover scenarios.
(5) *Space*. Space (also referred to as satellite) networks typically have long round-trip times which affect the transport-layer behavior.
(6) *Wireless-General*. This category catches research in which a generic wireless environment is used to motivate the existence of effects such as bit errors, but no coupling to any specific technology is provided.
(7) *Wireless-Unspecified*. This category reflects the case where the research considers a wireless environment but this is not explicitly modeled in the publication, so that wireless-related effects are not accounted for. Specifically, this is the case for some research concerning handover.

Wired domain.

(8) *Managed*. This category captures network environments in which a high degree of control exists over the network for the entire end-to-end path. This allows for appropriate dimensioning, traffic engineering, and quality of service (QoS) mechanisms to provide the desired network characteristics. Operator-owned IP-based signaling networks are a typical example of managed networks and an original design target for SCTP.

(9) *QoS*. This network environment refers to cases in which some QoS enhancing mechanisms are available for use in conjunction with SCTP. It covers environments which provide a better service than pure best effort but where complete control of the end-to-end path is not available.

(10) *Best Effort*. The best-effort network environment provides only best-effort end-to-end packet transfer service (e.g., best-effort Internet) without any assumptions about or restrictions on delay bounds, loss rates, etc.

3.4. Dimension 4: Study Approach

The fourth dimension captures the method used to obtain the results. Different methodological approaches have different benefits and drawbacks, and the results are strengthened if multiple approaches are used. This dimension uses a quite generic classification into the following five categories.

(1) *Conceptual*. The conceptual description approach describes and reasons about ideas, mechanisms, and functionalities in a general way without providing quantitative data to analyze the performance.

(2) *Analytical*. Analytical modeling tries to describe the essential behavior of an entity (such as a protocol) with a mathematical expression that, given some input parameters, provides some metric of interest. When a suitable expression has been derived, it can then easily be used to predict the performance of the entity under a range of conditions. However, in order to create a tractable formula, the expression must often be simplified, which introduces inaccuracies and highlights the need for model verification.

(3) *Simulation*. The simulation approach also uses abstract representations of the entity under study, but in this case, the abstractions are much more detailed and can include all relevant protocol functionality. Also, with simulations, there is a need to verify that the abstraction used in the simulation is correct and representative. Simulation allows a large parameter space to be explored and can provide considerable detail in the output.

(4) *Emulation*. In contrast to analytical modeling and simulation, the emulation approach uses actual protocol implementations running on real hardware. The emulation approach naturally captures the behavior of a protocol *implementation* and also allows for factors such as possible interaction effects with the operating system, device drivers, and communications hardware. Here, it is the behavior of the end-to-end connectivity that is abstracted to some degree by the employed emulator.

(5) *Live*. The live network approach entails performing experiments on a running communications network. This naturally includes all aspects, both at the endpoints and at the network level. However, live experiments are typically hard to control and repeat. Getting access to testbeds or live networks to the extent necessary may also be problematic in some instances.

4. SCTP RESEARCH ANALYSIS

The described taxonomy was applied to classify SCTP-related research. We proceed with the discussion of taxonomy results, presenting relevant cuts for each dimension in Section 4.1 and annual time series analysis in Section 4.2.

4.1. SCTP Research Profile

Figure 4 presents an overview of the investigated research articles for each of the taxonomy dimensions. Note that research in an article can cover more than one category within a dimension, and in that case, one article is assigned to multiple categories. Since the figure shows the percentage of articles that has research covering a particular category, the sum of all category percentages for a dimension may thus exceed 100%. To reflect the extent of this multiple assignment, we define the multiple categories index (MCI) as the ratio of the number of all category assignments for each dimension to the total number of articles classified.

Considering Dimension 1, protocol feature examined, Figure 4(a) clearly shows that the two major new SCTP features multihoming and multistreaming have attracted a dominating part of the SCTP-related research. Aspects of the multihoming feature are investigated in 56% of all articles, and the multistreaming feature is examined in 18% of the articles. Given that some articles address both features, the fraction of articles addressing these two major new SCTP features becomes 71%. A considerable number of articles are also related to the DAR extension (about 20% of all articles). However, most of the articles investigate DAR in conjunction with the handover aspect of multihoming. This is the main explanation for the relatively high value of the MCI ratio (1.39) for this dimension. Out of the remaining protocol features, the most explored is singlehomed congestion, a topic that has already been well studied for TCP. Relatively little attention in the research community has been spent on security issues and on the PR-SCTP extension (most of the time PR-SCTP appears in conjunction with the multistreaming feature), both covered by less than 8% of all articles.

For Dimension 2, application area, Figure 4(b) shows that nearly 45% of all articles use the bulk transfer application model. Interestingly, almost one third of all the articles classified as bulk transfer do not have any explicit specification regarding the application model that was used but had implicit indications towards the bulk model. The original application area of SCTP, signaling transport, is investigated by slightly more than 15% of all articles. Roughly the same number of articles has an unspecified application model, a quite common case especially within research devoted to transport-layer mobility. The multimedia category, reaching almost 18% of the articles, appears as a promising research direction for SCTP. The remaining categories account for less than 5% of the articles each, with the Web transfer and MPI categories having the highest numbers. In this dimension, most of the articles are assigned to only one category, leading to an MCI of 1.03.

Dimension 3, network environment (Figure 4(c)), also has very few multicategory articles (MCI is 1.04). Slightly more than a half of all classified articles (about 59%) are related to wireless environments. In contrast, the biggest single category belongs to the wired domain in which a best-effort network environment is used in about one third of all the articles. Research relating to the remaining two wired categories is limited, with either category examined by less than 10% of the articles. Heterogeneous networks (18%) is the most common wireless network environment, again, mostly because of the transport layer handover research. General wireless environments are investigated by a fair number of articles (12%), most of which are related to conceptual models or general handover schemes without any particular network clearly specified. While articles examining WLAN and MANET/VANET are not so numerous, they have

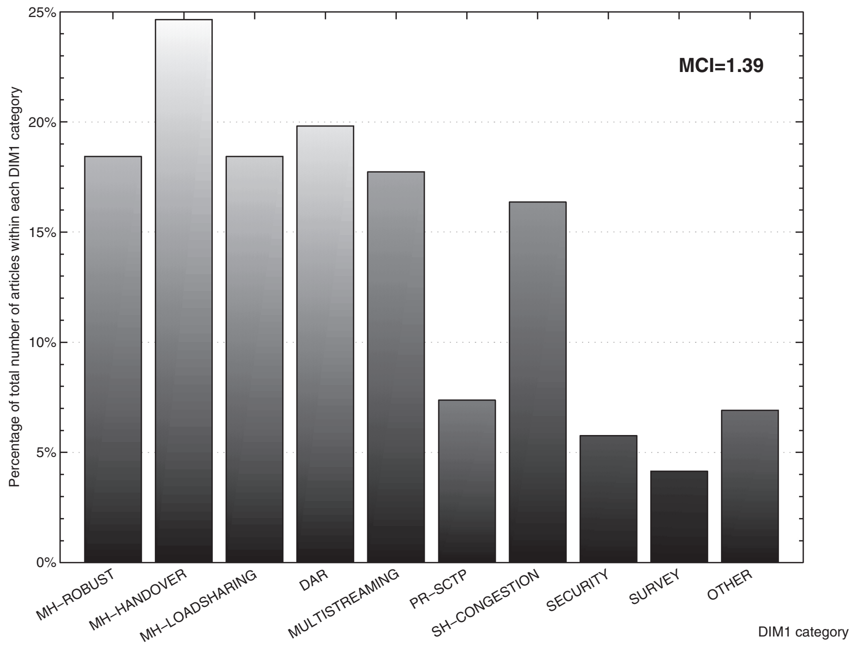

(a) Dimension 1: Protocol feature examined.

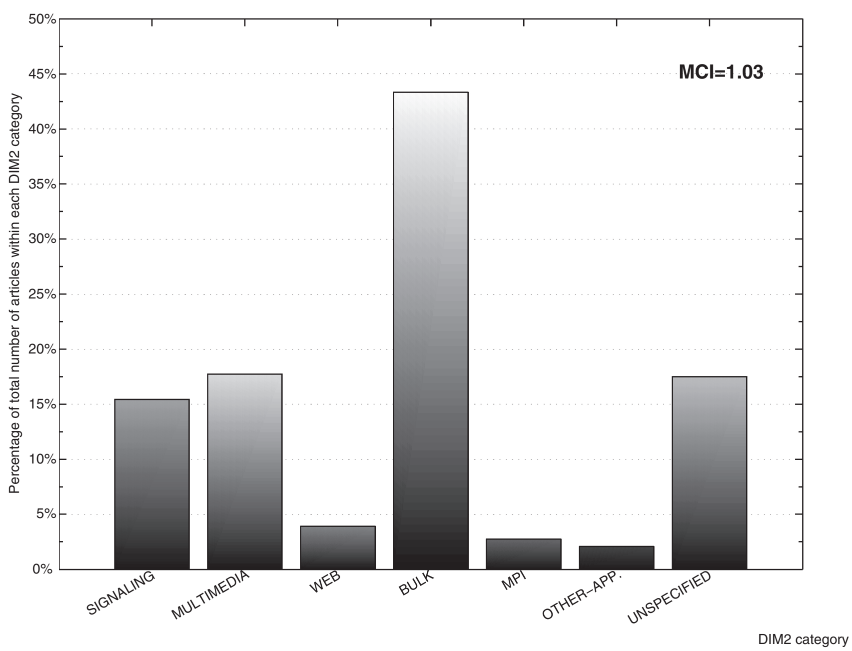

(b) Dimension 2: Application area.

Fig. 4. The percentage of articles classified within each category. The sum may exceed 100%, since an individual article can be classified into more than one category per dimension.

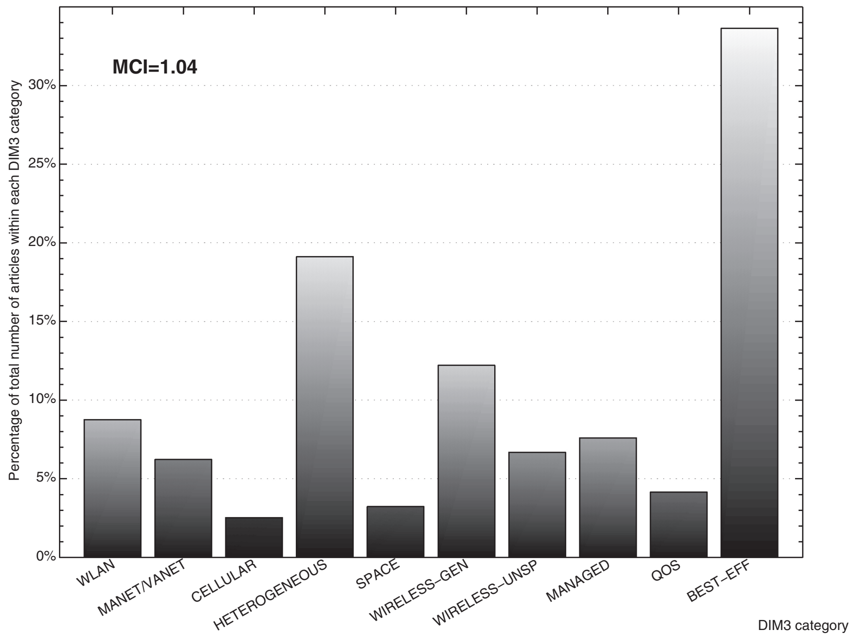

(c) Dimension 3: Network environment.

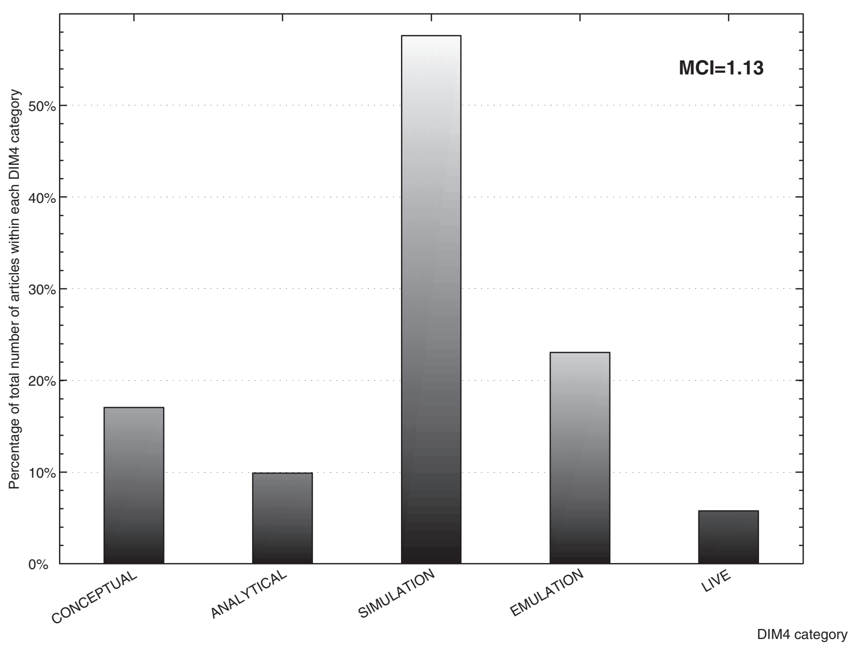

(d) Dimension 4: Study approach.

Fig. 4. (Continued).

mostly appeared during the last few years, and if this trend continues, these categories may increase in significance. The remaining wireless categories are examined by less than 10% of the articles.

Results for Dimension 4, study approach (Figure 4(d)), are influenced by some articles that combine two different study approaches for the investigation of a given issue (MCI is 1.13). This is especially the case for articles using analytical models which are verified by either simulation or emulation. More than a half of all SCTP-related articles are using simulation, typically either the ns-2[7] or the Qualnet[8] SCTP models which both are provided by the University of Delaware. Articles employing emulation to examine SCTP protocol implementations account for more than 20% of the total. Conceptual descriptions and analytical models are used to a lesser degree, both combined are used by slightly more than 25% of the articles, while less than 6% of the articles are involved with live experiments.

4.2. Category Annual Distribution

Figure 5 illustrates the annual distribution of the classified research. The graphics presented for each dimension shows the number of articles categorized into each category for each year, as well as the total number of articles published for that year. The difference between a particular bar and the corresponding line point illustrates the fact that some articles were classified into multiple categories within one dimension.

Dimension 1 (Figure 5(a)) provides an interesting illustration of the evolution of research related to the new features of SCTP. For multihoming, the initial research interest scoped on robustness, that is, the original application of the multihoming feature. Research related to the MH-Robust category presents a fairly stable annual contribution from the first year after the protocol specification was released, although research on robustness seems to have weakened somewhat over the last year. In contrast, the Dimension 1 category with the largest research interest, MH-Handover, represents a newer trend in SCTP research which started at the end of 2003. From that moment, we can observe an increasing number of MH-Handover articles over the last five years, with a peak of 23 articles in 2009. A similar trend can be seen for the third application of multihoming, loadsharing. Articles examining MH-Loadsharing started to appear in 2003, with 2004 already being the year with the largest number of articles, and a lower but fairly stable research effort since then. As mentioned before in Section 4.1, the DAR extension is strongly related to the MH-Handover feature and therefore follows its distribution. Articles examining multistreaming, the second novel feature of SCTP, provide a stable contribution over all inspected years, with the peak observed between 2006 and 2008. Similarly to MH-Robust, articles in the multistreaming category have a slight decline in the very last year analyzed. Singlehomed congestion, after being investigated by a considerable fraction of the articles in the first few years, lost in relative importance, although a notable increase has occurred in the last year with the highest number of articles occurring in 2009. A decreasing research interest has also affected survey contributions, and once SCTP became a fairly known protocol (about 2005), this type of publication appears less frequently.

For Dimension 2 (Figure 5(b)), articles examining bulk traffic applications have been quite clearly dominating this dimension over the last decade. Research on signaling and multimedia applications also shows stable trends over the analyzed years, with the latter category clearly taking over in the last two years (mainly due to VoIP applications of SCTP). For the remaining categories, it is harder to evaluate the tendencies, since with relatively little research already done, a single article makes a significant

[7] `http://pel.cis.udel.edu/`.

[8] `http://degas.cis.udel.edu/SCTP/`.

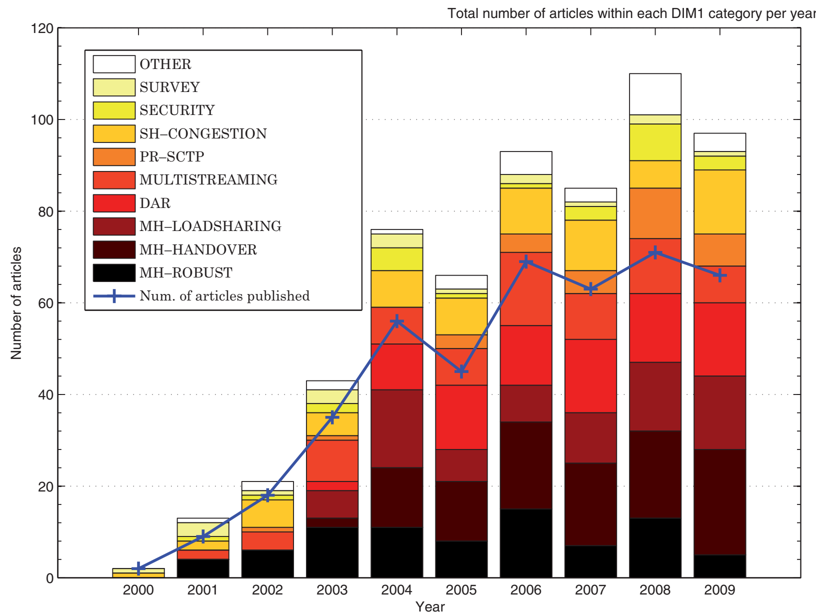

(a) Dimension 1: Protocol feature examined.

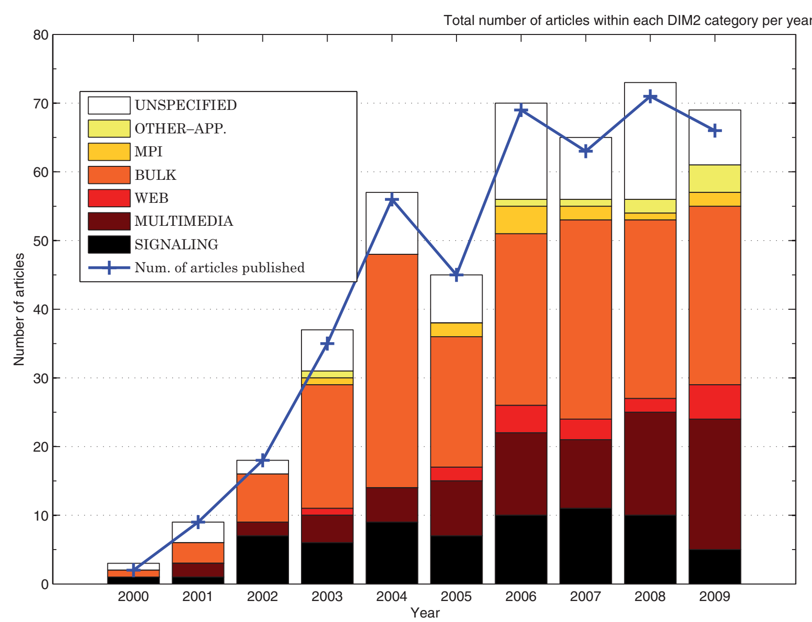

(b) Dimension 2: Application area.

Fig. 5. Annual distribution of all articles within each dimension.

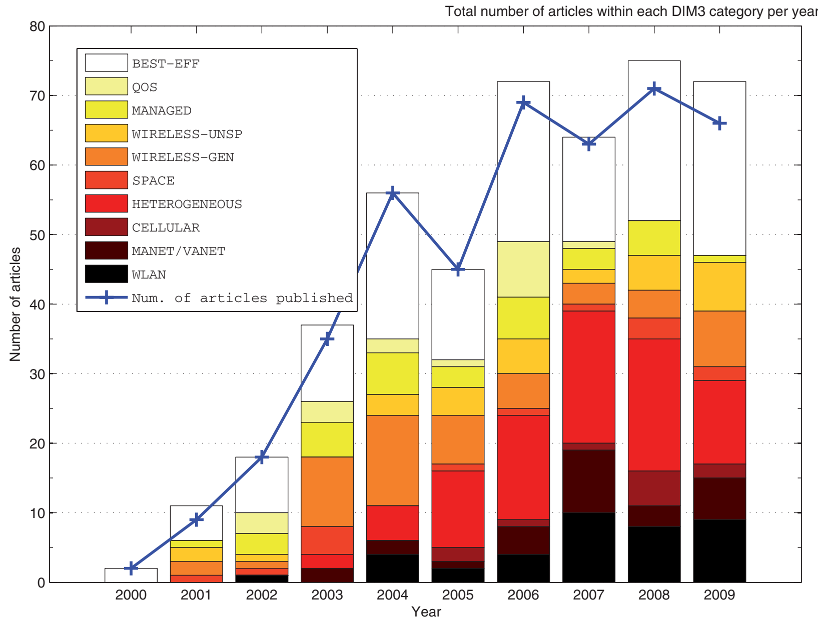

(c) Dimension 3: Network environment.

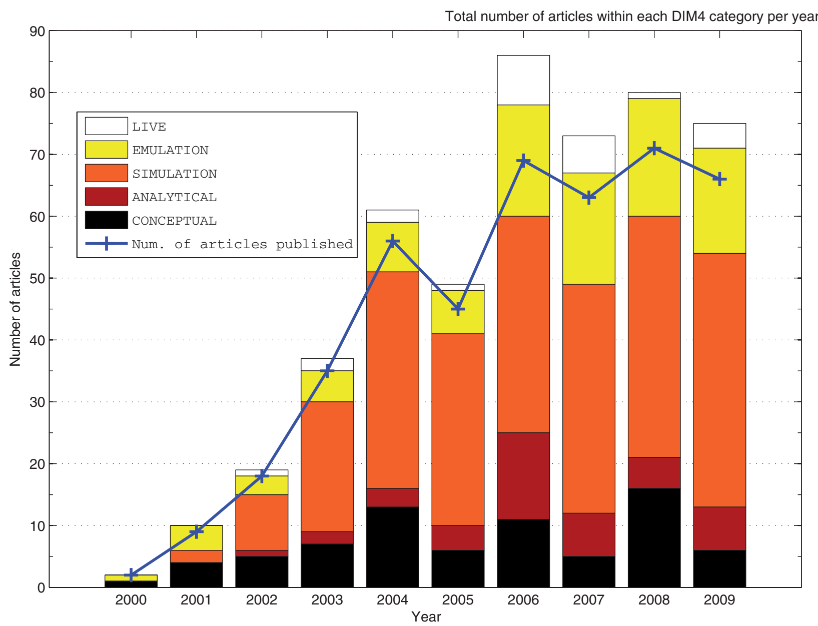

(d) Dimension 4: Study approach.

Fig. 5. (Continued).

difference. For example, articles considering Web applications over SCTP and SCTP for MPI applications first occurred in 2003; however, a consistent interest in both cases occur from 2005 and onwards. Research on other applications of SCTP mainly appeared in the last two analyzed years, but due to the small number of papers devoted to such applications, no stable trend can be inferred.

Dimension 3 (Figure 5(c)) is also dominated by articles examining one particular network environment—in this case, the best effort category of the wired domain. The first article examining the most common wireless category, heterogeneous networks, appeared in 2003. However, observing the evolution, it can be said that examining heterogeneous networks is one of the strong trends in recent SCTP research. In contrast, the considerable research devoted to satellite environments in the initial years radically declined after 2004. Another tendency that can be observed is that wireless research in recent years is less general, with the relative importance of the wireless-general category being decreased for the last four years. In contrast, for the WLAN and MANET/VANET categories, most of the publications fall within the last four years. Considering the relationship between the wired and wireless domains, it can be observed that while articles examining wired environments were in the majority before 2005, the last three years show a clear domination of wireless-related research.

In Dimension 4 (Figure 5(d)), the dominance of research based on the simulation approach is evident, starting from 2002 when the protocol model for the ns-2 simulator was released. The emulation approach, which has been used since the time when the protocol was just introduced, is used by a smaller number of articles. However, the number of articles using emulation shows a noticeably increasing trend over the last four years. Over the analyzed decade, research using conceptual and analytical approaches has also provided stable contributions (with peaks dating back to 2006 and 2008, respectively). Live experiments are used only by a small fraction of the studies, with the highest usage during 2006 and 2007.

5. SURVEY OF SELECTED EXAMPLES

As discussed earlier, multihoming is one of the key new features introduced in SCTP. It is also the feature that has attracted the most attention from the research community (more than half of all the analyzed articles). In this section, we elaborate on this feature and provide a survey of multihoming research, using our proposed taxonomy as a basis for the discussion. In line with the taxonomy subdivision of multihoming into the three aspects of robustness, handover, and load balancing, these areas are covered in separate sections.

5.1. Robustness

As mentioned earlier, the robustness aspect of multihoming was one of the major motivators for SCTP. At the time of design, the driving application area behind SCTP was the transport of signaling information in IP networks. In telephony networks, SS7 signaling is transported in dedicated SS7 networks with robustness built into the network structure at multiple levels. In contrast, IP networks provide only a best-effort service which does not match the requirements of signaling transport. The SCTP multihoming feature aims to ameliorate the decrease in robustness that results from migrating to a more generic and cheaper IP-based infrastructure from a more expensive dedicated SS7 network infrastructure. The research on the robustness aspect of multihoming can be divided into several subtopics. The most important subtopics concern retransmission-related issues that occur in a multihoming environment, failover detection and failover handling, and best path selection. Each of these topics is discussed next.

5.1.1. Retransmission-Related Issues. The SCTP retransmission behavior is decidedly different in a multihoming environment as opposed to a singlehomed. With multihoming, there are two or more potential paths that could be used to transmit the retransmission. The main cause for retransmissions are lost packets. Packet loss on the primary path typically implies that there is either temporary path overutilization or path failure on the primary path. In either case, sending retransmissions on the alternate path appears to be appropriate, since they then avoid the apparent problem on the primary path. This was the standardized behavior in the first RFC 2960 standard [Stewart et al. 2000].

The retransmission policy, that is, on which path the retransmissions should be sent and under which constraints, has been extensively studied by Caro et al. [2003a, 2003b, 2004b, 2006b]. Three basic retransmission policies were evaluated: send all retransmissions on the alternate path, send all retransmissions on the primary path, and a hybrid policy that sends fast retransmissions on the primary path and timeout retransmissions on the alternate path. The earlier studies [Caro et al. 2003a, 2003b, 2004b] use a symmetric configuration in which both paths have the same bandwidth and delay characteristics, whereas the later study [Caro et al. 2006b] considers also the case of path asymmetry. After their comprehensive evaluation of both failure and non-failure scenarios, the recommendation of Caro et al. is to adopt the hybrid retransmission policy that sends fast retransmissions to the same IP address as the original transmission and sends timeout retransmissions to an alternate peer IP address. None of the evaluated policies is the best in all scenarios, but the recommended policy is the most robust and shows good performance for all evaluated scenarios. This new hybrid retransmission policy has been proposed as a recommended update, first in RFC 4460 [Stewart et al. 2006], and then introduced into the SCTP specification in RFC 4960 [Stewart 2007].

In a multihoming environment, there is also the question of which path to use when sending SACKs from the receiver to the sender. If the alternate path has considerably lower delay than the primary path, it could be beneficial to adapt the protocol behavior, and the earliest research looking into this issue is by Jungmaier et al. [2001]. In their work, the authors consider a satellite scenario with a high-bandwidth, high-delay satellite primary link and a low-bandwidth, low-delay terrestrial backup link. A modification is proposed to the SCTP retransmission algorithm so that all SACK chunks that contain gap reports, and thus possibly indicate losses, are sent via the low-delay backup path. In the studied scenario, this results in a reduction in the packet delivery time of over 30%. Additional simulations supporting this idea are also later provided by Jungmaier and Rathgeb [2006] in which the improvement in throughput is examined for a similar satellite scenario but now also with bit errors. As the bit error rate (BER) of the satellite link increases, the proposed optimization provides increasing benefits leading to almost a doubling of the throughput as the BER reaches 2×10^{-6}.

Qiao et al. [2007] also consider an asymmetric setting in which the delay for the secondary path is lower than for the primary path. They do not consider the issue of which path to send the returning SACKs on, but rather on the problems created by disordered SACKs and fast retransmit behavior. Four different fast retransmit policies are examined, considering the impact of the Max.Burst parameter introduced in RFC 4460 [Stewart et al. 2006] in order to limit the maximum number of packets that can be sent out at the same time. In their study, no single policy is found to be the best for all evaluated configurations, but when the bandwidth is high, the use of the Max.Burst parameter appears to have a positive effect.

5.1.2. Failover Detection and Failover Handling. The goal for the failover detection mechanism is to quickly and reliably detect when the primary path has failed so that traffic instead can be sent on a backup path. To detect transmission problems, SCTP keeps track of missing acknowledgments at the sender. To distinguish a path failure from temporary congestion, the sender maintains an error counter which counts the number of consecutively occurring timeouts. If the error counter of the primary path reaches a set threshold, *path max retransmit (PMR)*, the primary path is considered unavailable or unreachable, and a failover is performed. If, on the other hand, a SACK is received for data sent on the primary path, the error counter is reset to zero, and the previous timeouts are assumed to be caused by temporary congestion. The tuning of the PMR value is thus a critical factor for failover performance, as too large a value will delay the detection of path failure, whereas too small a value may lead to spurious (or unnecessary) failovers.

Early work on failover detection was performed by Jungmaier et al. [2002]. This research examines failover performance in an SS7 application environment and comes to the conclusion that the recommended default values (as specified in RFC 4960 and kept unchanged since the initial protocol specification) for several SCTP parameters related to failover will give unacceptable performance for signaling applications. This problem occurs as the default parameters are set for general Internet use and not tuned for signaling environments. Simulations examining the effect of lowering PMR as well as the minimum and maximum bounds for the retransmission timeout (RTO.min and RTO.max) show that the requirements for SS7 signaling transport can be fulfilled by adjusting these values appropriately. The inappropriateness of the default SCTP failover parameter settings in a signaling environment and the need for parameter tuning is also confirmed by emulation experiments performed by Grinnemo and Brunstrom [2004]. The results from their study show that a PMR of 2 in combination with a lowered RTO.min results in a failover time below 2 s in all studied scenarios—a performance sufficient for most SS7 applications. Later experiments by Grinnemo and Brunstrom [2005] also examine the effects of bursty cross traffic and router buffer queue sizes in an SS7 context. The results from their experiments show that in the presence of cross traffic, large router buffers lead to highly varying delays, which in turn inflate the RTO calculation at the sender and increase the failover time. The results thus highlight the importance of having relatively small router buffers. The study by Eklund et al. [2008] illustrates the effect of the SACK delay on failover performance. For traffic consisting of individual signaling messages at low intensity, the default SACK delay of 200 ms more than doubles the failover time as compared to no SACK delay, whereas the impact of the SACK delay for high intensity or bursty signaling traffic is limited. Building on several of the works just mentioned, a coherent treatment of how to configure the SCTP failover detection mechanism for carrier-grade telephony signaling, including practically usable configuration recommendations, is provided by Eklund et al. [2010]. The authors also suggest relaxing the exponential backoff that forgoes a retransmission timeout in SCTP as an alternate or complementary way of optimizing the failover detection.

The positive effect of a lower PMR has also been shown in extensive simulation studies by Caro et al. [2004a, 2006a]. They do not specifically consider signaling environments, but rather a general Internet setting, and also use considerably larger path-propagation delays. Their results show that the failover performance can be increased by lowering the PMR all the way down to a value of 0. Even though a PMR of 0 may lead to a large number of spurious failovers, it still leads to improved goodput in the studied bulk transfer scenarios. Based on their results, the authors suggest a revised more aggressive failover mechanism in which transmission of new traffic is immediately migrated to a new active path as soon as a timeout occurs. The abandoned primary path is probed for reachability with heartbeats and not marked as inactive (or failed) until the PMR for the path is exceeded. If a probe is successful, the transmission of new traffic is migrated back to the original primary path. The idea is thus to continue data transmission on a different path during the failure detection period.

As reported by Noonan et al. [2006], the failover mechanism is susceptible to stall when some unusual failover conditions occur, such as an erroneously low estimate of the RTO value for the backup path or network errors that affect only SACKs. Noonan et al. propose several preventive measures. The most general solution suggested aims to disambiguate the SACK information for data chunks that, due to retransmission, have been transmitted over both the primary and the backup paths. This is done by including a tag in the SACK that indicates on which path the acknowledged data chunk was received. The SCTP failover mechanism can also be combined with an application-layer failover mechanism. This has been proposed by Yoo et al. [2002] for the case of the Diameter application in which SCTP handles failover between different network paths to the same authentication, authorization, and accounting (AAA) server, and the AAA client handles failover between different AAA servers.

5.1.3. Best Path Selection. Closely related to failover detection is best path selection that dynamically selects the most appropriate path for data transmission. The basic idea of best path selection is to change path not when the primary path has failed, but rather to change to another path when the other path provides “better” transmission conditions. The strict primary and backup semantics of the paths are thus relaxed to achieve potential performance benefits. In a way, best path selection can be viewed as making an immediate but conditional failover at the first sign of path problems (e.g., a retransmission trigger or increased transmission delay). In contrast to the regular failover mechanism, best path selection schemes may also have additional functionality for probing the paths and selecting which path to use. The decision on which path to use for data transmission is made by the sender, but this choice may be directly dependent on information provided by the receiver.

The key issues for any best path selection scheme are thus to determine how and when the transmission conditions of the different paths should be compared. In Noonan et al. [2004] the selection of the path to use is based on either a utility function or a utility-cost function. The function is defined by the application and can thus be tailored to the requirements of each application. A hypothetical media application is considered. The delay and jitter experienced over all active paths are monitored through data traffic and heartbeats, and the utility function is based on a combination of the delay and jitter over the path. The utilities of the different paths are compared at regular intervals—every 9 s (which corresponds to every 30 packets for the evaluated application)—and the path with the highest utility is selected for the next period. When the utility-cost function is used, the utility of each path is weighed against the cost of sending packets over this path. The path selection is controlled by the receiving application and conveyed to the sender by transmitting a special control chunk. According to the presented experimental results, the use of the utility-cost function leads to an 18% utility gain in comparison to the best static selection if the background traffic load varies dynamically over both paths but may lead to a slight decrease in utility when the background traffic load is stable.

Similarly, the use of a profiling framework that allows customized profilers to infer different path characteristics from heartbeat information is suggested by Gauch and Nishida [2006]. Different applications can thus use different profilers to control path selection according to their requirements. Their prototype implementation in FreeBSD includes two sample profilers. One profiler selects the path with the lowest loss probability, performing a comparison between the paths after each packet loss. The other profiler selects the path with the lowest round trip time (RTT), performing a comparison between the paths when a significant increase in RTT is noticed on the current primary path. No experimental results on the use of the profilers are provided.

The use of bandwidth estimation techniques for selecting the path with the most available bandwidth on each timeout expiration event is suggested in the WiSE scheme [Fracchia et al. 2005, 2007]. In addition to the multihoming-related change, the scheme also proposes a Westwood-like congestion control where the estimated path bandwidth is used to help separate congestion-related losses from those caused by wireless errors. The multihoming change proposed by WiSE could be seen as a failover modification, where failover occurs for the first timeout, but only if the backup path at that instant has more available bandwidth than the primary. The bandwidth on the primary path is estimated based on the flow of returning SACKs, whereas the heartbeat mechanism is extended to send a six-packet train that is used to estimate the bandwidth on the alternate path. Simulation results indicate that WiSE is able to select the less congested path which also leads to an increase in throughput. A closely related approach is presented by Casetti et al. [2006, 2008] in a scheme named AISLE. This approach is discussed in the context of multiple overlapping wireless access networks in which a host needs to automatically and autonomously select the access network that delivers the best performance. Again, bandwidth estimation is used to select the path with the most available bandwidth. Unlike WiSE, AISLE triggers a potential path change also on fast retransmits and not only on timeouts. In order to avoid frequent oscillations, AISLE has a time hysteresis built in that ensures that two path changes are separated by at least a minimum time period (60 s is used). Simulation results indicate that AISLE is able to distribute the load almost perfectly over the available networks, both in a scenario with multiple overlapping WLAN networks and in a mixed WLAN-UMTS scenario.

5.1.4. Classification of Related Articles. The distribution of all research articles on multihoming-robustness over the other dimensions in our taxonomy is shown in a scatter plot in Figure 6. The scatter plot provides a visual representation of which combinations of application area (Dimension 2), network environment (Dimension 3), and study approach (Dimension 4) were considered in the articles classified as MH-ROBUST in Dimension 1. The application area is shown on the y-axis and the network environment on the x-axis. Each article is placed at the intersection of the appropriate lines and represented by a symbol that corresponds to the used study approach. A small random scatter value is added to separate the entries around the intersection to avoid total overlapping. As can be seen in Figure 6, the articles form two main clusters around signaling traffic in managed networks and bulk traffic in best-effort networks. As multihoming was introduced to provide the robustness expected by signaling applications, it is perhaps a bit surprising to find that the signaling area is not the largest application area. Instead, it is bulk transfer that is the most represented application area. One potential explanation for this is that research on TCP, to a large extent, has focused on examining bulk transfer and that this has carried over to SCTP research. Although research performed with a focus on bulk transfer may be applicable in a signaling context, the focus of signaling transport is on minimizing the message transfer delays of the individual signaling messages rather than on maximizing the throughput of large data transfers, which is the main issue for bulk transfer. Consistent with bulk transfer being the dominant application area, best effort is the most studied network environment. Simulation is used by a clear majority of the studies, although for research on signaling in managed networks, it can be seen that emulation is the most popular approach.

5.2. Transport Layer Mobility

Using multihoming only for robustness-related purposes can be seen as not taking full advantage of the benefits multihoming can offer. Therefore, the idea of using the

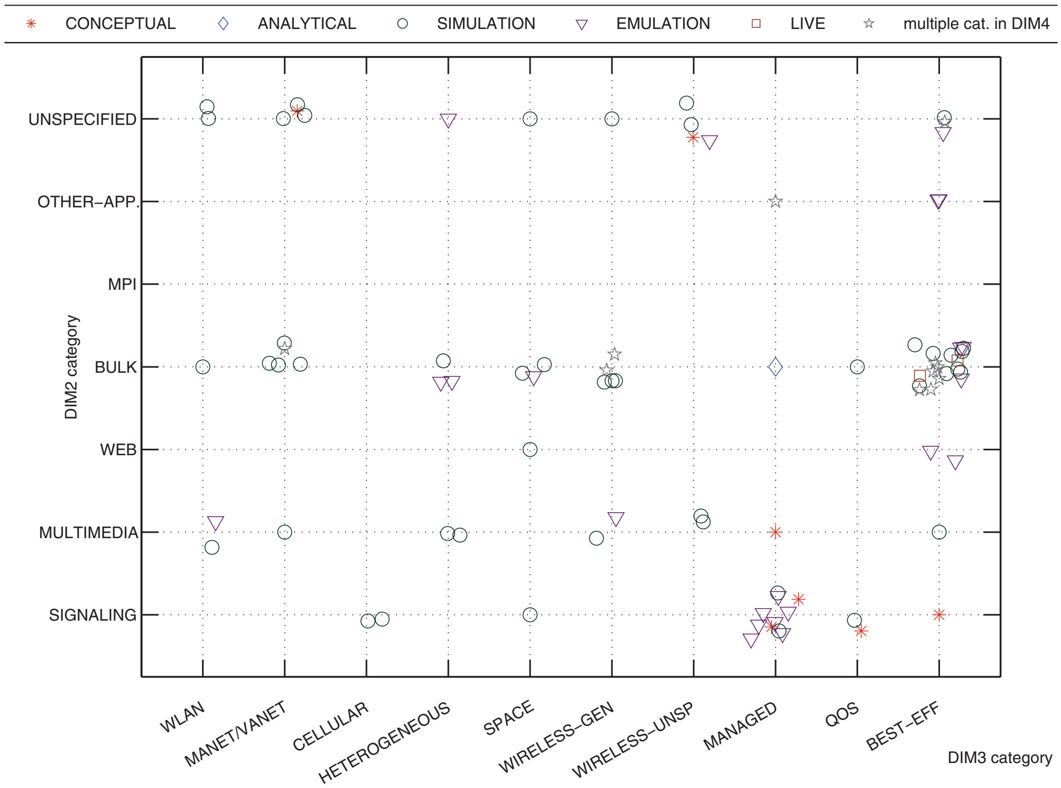

Fig. 6. Scatter plot of all robustness-related articles.

multihoming feature for transport layer mobility emerged in late 2003, as soon as the work on the specification of the DAR extension (providing means for dynamic address manipulation within an established association, as described in Section 2.2) got to an advanced stage. The key idea behind transport layer mobility is to handle mobility on an end-to-end basis. This removes the need for extending the network infrastructure with specialized mobility support to allow user sessions to roam uninterruptedly between IP subnets. Under such a view, the DAR extension allows SCTP associations to survive possible IP address change(s) when *handing over* between different networks or between subnets within a network. Consequently, both protocol features, multihoming and DAR, are usually put together as the two main enablers for handover support within an SCTP-based transport layer mobility solution. In this context, standard SCTP together with the DAR extension is often referred to as mobile SCTP (mSCTP) [Riegel and Tuexen 2007]. Adding support for handover at the transport layer has important advantages over traditional network layer approaches (e.g., Mobile IP and its derivatives), such as the potential to perform smooth handovers (by exploiting multihoming in the transition process), scalability (handover functionality is distributed within end nodes), efficiency (encapsulation and tunneling overhead is avoided), and fault tolerance (solution is not dependent on specific centralized network functionalities). Moreover, adding handover support at the transport layer can be easily complemented with location management solutions (e.g., Dynamic DNS (DDNS) or network-layer or application-layer registrars) in order to have a complete mobility solution (i.e., handover and location management).

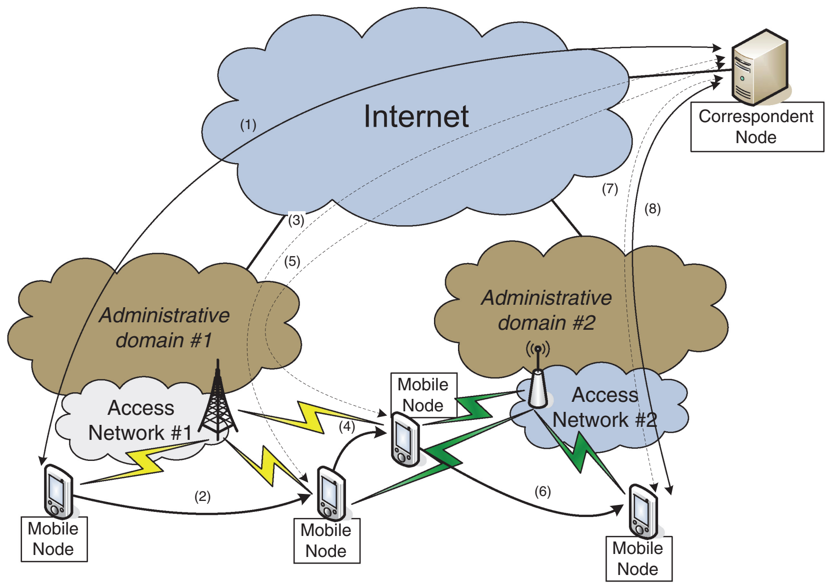

Fig. 7. mSCTP architecture and operations.

Exploitation of multihoming within SCTP-based transport layer handover solutions means that mobile devices must be able to keep at least two connections simultaneously (i.e., multihomed terminals). This is a common case within multimode terminals of different technologies, such as devices with cellular and WLAN network interfaces. However, it is worth noting that SCTP-based transport layer mobility solutions are not strictly coupled to multihoming support. In this regard, work by Honda et al. [2007] proposed an SCTP-based handover solution for singlehomed nodes that only requires the DAR feature. In any case, this work of Honda is the only exception to all other analyzed handover-related articles which assume multihoming support within the mobile device. A categorization of different mobility scenarios, including different single- and multihomed configurations where SCTP-based solutions can be applied, is provided by Budzisz et al. [2008]. The conclusion of this work is that the most common mobility scenario in which to exploit mSCTP is a heterogeneous wireless network scenario with a multihomed mobile node (MN) and singlehomed correspondent node (CN). Figure 7 illustrates an example of an mSCTP-based handover process for such a typical handover scenario in which a multihomed MN moves from one network to another while maintaining communication with a singlehomed CN. When establishing an mSCTP association between MN and CN, both nodes first exchange the lists of IP addresses valid for the communication (1). Only one of the destination addresses is selected to send the data to (primary path), whereas all remaining destination addresses serve only for backup purposes. As long as the MN stays in the area where the initially defined IP addresses are available, there are no mobility-related concerns. If the MN moves to an area where a new IP address (an address that was not included in the initial list) has to be used by the MN (2), then as soon as the new IP address is known at the MN side, it can be communicated to the CN using specific control messages (address configuration (ASCONF) chunks defined in the DAR extension) (3). Modifying the IP address(es) of the association increases the risk of association hijacking, as described in RFC 5062 [Stewart et al. 2007]. Therefore, to prevent hijacking attacks, the ASCONF chunk must be sent in an authenticated way (a special AUTH chunk is bundled before the ASCONF chunk), as described in RFC 4895 [Tuexen et al. 2007]. Once the new IP address has been added to the association, the MN may decide if it should still send data to the old IP address or in an appropriate moment (4), subject to handover policy, switch the primary path to the new IP address (5). After the address change, the transmission can continue uninterrupted on the new IP address (a smooth or even seamless handover). When leaving the old access network (6), the unnecessary IP address(es) can be removed (7), as further transmission goes on (8).

In the next section, we introduce the initial work conducted on mSCTP along with some studies addressing performance comparisons of SCTP-based schemes to other mobility solutions. With that as a base, the remaining sections cover some of the most relevant and challenging issues addressed so far within this research topic: leveraging failover schemes to manage handover, improving handover decisions by adding support from other layers, and enhancing data transmissions during the handover transition process.

5.2.1. Initial Work and Comparisons to Other Mobility Schemes. Most of the initial work on mSCTP focused on conceptual descriptions of the handover mechanism and identification of its main open points. One of the examples of such conceptual works is the article by Koh et al. [2004] in which the basics of the mSCTP operation are described along with an initial performance analysis of the solution in a WLAN scenario. Link layer signal strength information is used to configure the triggering rules for adding a new IP address to and changing the primary IP address of an ongoing SCTP association. The obtained results corroborate that the envisaged concept can constitute a competitive approach when compared to traditional network layer solutions. Another of the initial works on mSCTP by Ma et al. [2004] provides a comprehensive description of an mSCTP-based handover scheme for heterogeneous networks, focusing in particular on UMTS-to-WLAN and WLAN-to-UMTS handovers. A detailed evaluation of such handover procedures is provided using the overall handover delay as a metric. It is shown that an mSCTP-based handover scheme can be successfully applied in the envisaged scenario and that the same solution can be practically extended to support vertical handover in any type of heterogeneous wireless environment.

Further evaluation of mSCTP in terms of handover delay, signaling cost, dropping probability, and overall throughput in a homogeneous WLAN scenario is given by Argyriou and Madisetti [2007]. Argyriou and Madisetti relate the obtained results to the MIP and Hierarchical Mobile IP (HMIP) schemes. For the simplest handover policy involving change of path as soon as the new IP address is operational, the mSCTP scheme is capable of achieving results comparable to HMIP (being less than 5% worse than HMIP in terms of throughput) while providing a more scalable solution, where no specialized mobility support is needed. This work is complemented by an empirical evaluation of mSCTP in comparison to MIP and SIP schemes performed by Zeadally and Siddiqui [2007]. Zeadally and Siddiqui conclude that although mSCTP can considerably outperform both schemes in terms of handover delay and throughput in all tested scenarios (e.g., 31% lower latency if compared to SIP and 55% lower latency when compared to Mobile IP for WLAN to Ethernet handover), an important shortcoming is the inability to operate in networks that use network address translation (NAT). In such a network (e.g., a network with dynamic and private IP addressing), NAT assigns a new port number that causes the SCTP association to be dropped. Readers interested in more details on SCTP support for NATs should refer to Tuexen et al. [2008] and Hayes et al. [2009].

5.2.2. Reusing Failover Mechanism. Apparently the most crucial challenge in mSCTP-related research is providing an optimal path management. The simplest approach among the investigated handover strategy proposals or when the handover triggering mechanism is not explicitly specified is to reuse the standard SCTP failover to trigger the path change within the handover process. Failover in such a handover context is specifically evaluated by Budzisz et al. [2006, 2008] who suggest decreasing the default PMR and RTO.min values of standard SCTP in order to adjust the failover mechanism to handover needs. Nevertheless, Budzisz et al. conclude that the standard SCTP failover-based mechanism is completely unsuitable for real-time applications; however, they do not completely rule it out for non-real-time applications. Similarly, the suitability of the failover mechanism for handling handover is evaluated by Noonan et al. [2006] who propose an additional improvement—a so called association routing table (ART) at the transport layer. ARTs are created in such a way that each destination has assigned a different source address (if possible). ARTs are synchronized at both multihomed endpoints in order to improve switching efficiency in the presence of network failures. Performance of the SCTP failover mechanism in a handover context has also been extensively evaluated in experimental setups. As an example, the evaluation in an IPv6 UMTS-WLAN testbed shown by Bokor et al. [2009] further confirms the need for adjustments if the failover mechanism is to achieve handover targets.

Recall that several mechanisms for enhanced failover detection and best path selection were discussed in Section 5.1. It is important to stress that schemes like WiSE [Fracchia et al. 2005, 2007] or AISLE [Casetti et al. 2006] can also serve as valid handover solutions.

5.2.3. Improving Handover Decisions. Aiming at improving the performance of the basic failover scheme presented in Section 5.2.2, lower-layer support could be considered in the handover decision-making process. The handover problem becomes especially complex in heterogeneous scenarios in which information from various networks regarding different link features, such as available bandwidth, security, monetary cost, as well as end-user preferences, should be simultaneously taken into account to make a successful handover decision. Events and parameters from lower layers (in different units) must be translated to the unified values of a general handover cost function in order to provide hints and triggers for the transport layer, for example, RFC 4957 [Krishnan et al. 2007] provides some patterns for triggers for GPRS, CDMA2000, and IEEE 802.11 link layers.

In the case of mSCTP, only quite simple handover policies have been analyzed so far. Most typically, the considered handover policies have been based on the relative signal strength criterion together with a form of hysteresis, usually with an aggressive and a conservative threshold value, to trigger appropriate DAR control messages. As an early example, Chang et al. [2004] extended the proposal introduced by Koh et al. [2004] (described in Section 5.2.1) to use link layer signal strength information to govern the address manipulation process and trigger the handover. Having the transport layer aware of the mobility thanks to the link layer indications, Chang et al. additionally introduced error and congestion control enhancements to better adapt mSCTP to the handover scenarios and reduce handover delay, losses, and loss recovery time. Another example of a scheme that reuses link layer information in the handover process is the proposal introduced by Kim et al. [2006]. The design is based on collecting link layer events (e.g., interface up/down) by the link layer monitoring module and processing them in the handover decision module at the transport layer, which finally triggers the handover in the most appropriate moment. Kim at al. demonstrate in a simple experimental setup on a Linux platform that their proposal is able to improve the throughput and decrease handover delay of the mSCTP handover scheme in heterogeneous (WLAN and cellular) network scenarios.

An alternative approach to handover triggering is presented by Chang et al. [2007] in which the handover strategy is based upon information on available wireless bandwidth calculated using link layer information in combination with contention probabilities obtained from periodically sent heartbeat probes. Such an approach is shown to improve throughput performance, for example, in the tested homogeneous WLAN scenario, even up to 60% over the regular mSCTP scheme; nevertheless, it raises serious doubts about its applicability in the case of heterogeneous wireless networks. In the context of lower-layer support enhancements, Budzisz et al. [2008] show that an inappropriately adjusted handover policy based on link layer information may result in a decrease of the performance, as compared to the previously discussed standard SCTP failover-based handover policy.

The idea of supporting handover decisions with cross-layer information can be extended to incorporate multiple layers, for example, see Fitzpatrick et al. [2009]. Fitzpatrick et al. present a scheme called Endpoint Centric Handover (ECHO), targeted for VoIP applications. ECHO incorporates several cross-layer metrics using noise and link quality information from the physical layer, availability of the access network from the network layer, and end-to-end characteristics from the transport layer. The ITU-T E-Model for voice quality assessment is used to map these metrics to a user-perceived mean opinion score (MOS) as a quality metric for VoIP. The final handover decision is made based on the MOS results for each candidate access network, resulting in improved overall VoIP performance.

5.2.4. Enhancements for an Ongoing Handover Process. An additional way of improving the performance of an mSCTP-based handover scheme is to introduce upgrades during the ongoing handover process. One of the early studies by Kashihara et al. [2004] proposed severe changes to SCTP, disabling its congestion control and data retransmission mechanisms and including changes to the path error accounting algorithm and the corresponding PMR limits and finally introducing the possibility of sending duplicate packets when more than one path is available. The changes were motivated by the fact that, originally, SCTP was optimized for non-real-time communication and that all mentioned changes aimed to adapt SCTP for real-time handover scenarios. The provided evaluation in a heterogeneous scenario (including WLAN and cellular networks) is somewhat insufficient to justify broader adoption for such considerable changes, as previously described. The idea of sending duplicate packets among simultaneously available paths has also been discussed in works by Aydin et al. [2003] and Aydin and Shen [2005] that introduce a scheme called cellular SCTP (cSCTP). cSCTP sets the congestion window on both the old and the newly obtained path to half of the value it had on the old path before the handover and starts sending duplicate packets. This scheme does not, however, provide any kind of estimation of available bandwidth on the newly obtained path before starting the transmission. The provided work mainly focuses on the conceptual description of the idea and lacks a broader evaluation that reflects the need for the proposed changes. Along with the idea of sending duplicate packets when multiple paths are available in a handover scenario, a proposal to use load balancing, first devised by Goff and Phatak [2004], has also been raised. Goff and Phatak argue that introducing load balancing to the transport layer mobility has the potential not only to increase the throughput but also leads to increased fairness among the users. In their initial experiments, Goff and Phatak use the multistreaming feature to facilitate loadsharing (for ease of the practical implementation) so that each stream is handled on a separate path. Practically, no further work followed in this direction until the loadsharing solution for SCTP was properly devised, as specified in more detail in Section 5.3. Huang et al. [2007] brought back this idea, presenting the design of a complete transport layer mobility scheme that takes into account loadsharing as a possible enhancement. Recently, the idea of using transport layer loadsharing in mobility scenarios has been further considered by Budzisz et al. [2009a, 2009b]

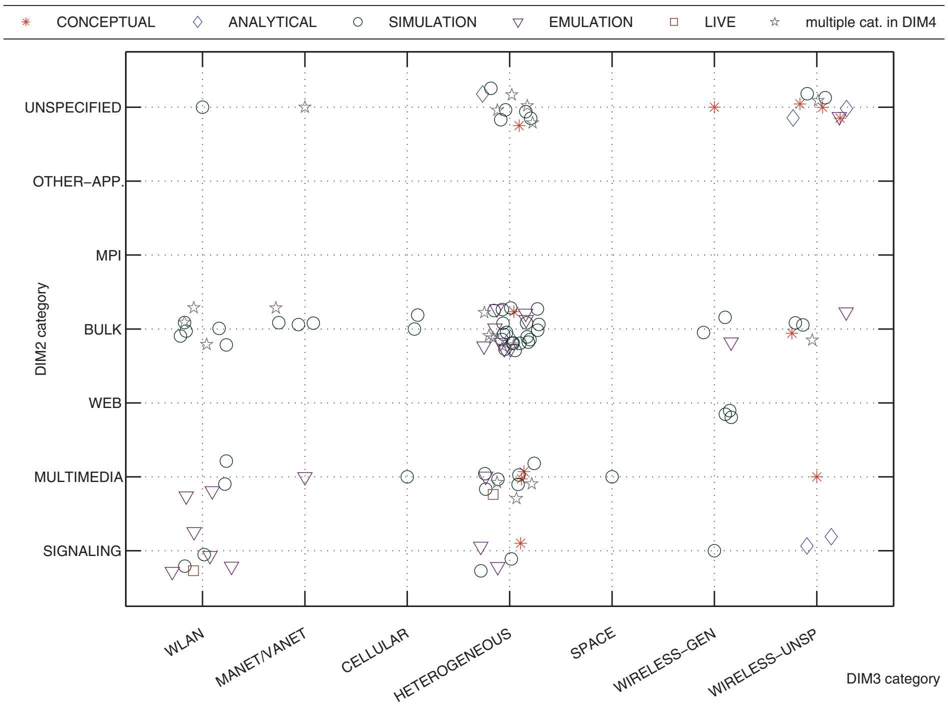

Fig. 8. Scatter plot of all handover-related articles.

in the design of a scheme called mSCTP-CMT-PF which reuses concurrent multipath transfer (CMT), the most common loadsharing solution for SCTP. In later work, Huang et al. [2009] also test a CMT-based loadsharing scheme.

Modification of the retransmission mechanism during an ongoing handover is also proposed by Ma et al. [2007]. They propose the so called smart fast retransmission mechanism that accounts for wireless channel and handover losses (i.e., losses that occur in packet rerouting processes while executing the handover), therefore reducing the risk of sending retransmitted packets to an already unavailable path. The results illustrate that this problem can be of particular concern, specially during WLAN to cellular forced vertical handovers.

5.2.5. Classification of Related Articles. Figure 8 illustrates the distribution of handover-related research over the remaining dimensions of the taxonomy. The first thing to notice is the unbalanced distribution with regards to the applications considered in the studies. If the application is explicitly specified, it is usually bulk transfer. Signaling and multimedia applications are less frequent, and such research cover more specific solutions influencing the design of the proposed handover scheme, for example, VoIP in the works by Fitzpatrick et al. [2006, 2009]. In contrast, among the network scenarios analyzed, there is more diversity with almost every wireless category being represented within the first five years since the first mobility-related article was published. The most typical handover scenarios analyzed include heterogeneous networks (usually WLAN and cellular) and wireless-general or unspecified wireless networks in case of conceptual works. mSCTP-based handover schemes are also evaluated in homogeneous networks, mainly in WLAN, and considerably less often in cellular networks. Future trends in handover-related research may also consider MANET environments that so far have had only a few contributions. Rather unsurprisingly, the vast majority of the research dedicated to transport layer mobility uses the ns-2 SCTP model to evaluate the proposed ideas. Only a few works go beyond this scheme and provide results from emulated environments, most of them using the Linux kernel SCTP (LK-SCTP) implementation.[9]

5.3. Loadsharing

Transport layer loadsharing is another application that extends the use of SCTP multihoming relative to what is defined within the standard protocol specification. Loadsharing techniques allow for data striping across multiple network interfaces and can be potentially applied at the network, transport, or application layer. Among such alternative approaches, transport layer loadsharing allows the application to be isolated from any modifications and, as compared to a network layer scheme, has a considerable potential to improve overall transmission performance and increase network efficiency. Decisions on simultaneously sent data are made on an end-to-end basis and thus it is easier to avoid unnecessary fast retransmissions or spurious timeouts while sending more data [Goff and Phatak 2004]. Despite its potential benefits, a challenging issue that a transport layer loadsharing solution based on SCTP has to face is packet reordering at the receiver due to simultaneous data transfer over multiple paths. Packet reordering could significantly deteriorate SCTP performance, since the congestion control algorithms in standard SCTP are derived from TCP and hence do not work well when reordering is common. As stated in Section 2.1, congestion control in standard SCTP is applied to the entire association. However, separate sets of congestion control variables (cwnd, sstresh, and partial_bytes_acked) are kept for each of the destination addresses of a multihomed peer. Thus, to successfully implement loadsharing within SCTP, the management of the send buffer and the congestion control must be updated to take into account the problems of sending data over multiple paths using a single sequence-number space and the consequences of sender-introduced reordering. So far, there is no commonly defined extension that facilitates loadsharing for SCTP. Therefore, the most important proposals will be examined here in detail in three separate subgroups: one subgroup covering initial work envisioned about loadsharing support in SCTP, a second one covering the most common loadsharing scheme so far, that is, CMT, and the last subgroup encompassing more sophisticated loadsharing proposals that make use of smarter packet-scheduling functionality at the sender side.

5.3.1. Initial Work. One of the first proposals for loadsharing with SCTP, called LS-SCTP, has been brought up by Abd el Al et al. [2004a; 2004b]. LS-SCTP separates flow control, handled per association, from congestion control that for loadsharing needs to be handled per path. Therefore Abd el Al et al. propose the introduction of two additional chunk types to carry data and related acknowledgments in LS-SCTP. Both chunks are backward compatible with their corresponding standard SCTP chunks. The only difference is the additional sequence numbers added to facilitate congestion control on a per-path basis, thus explicitly tracking the packets sent to a given destination. The proposed solution also offers a modified path-monitoring mechanism with more frequent heartbeat probing to avoid stalling the application on an inactive path. Despite that LS-SCTP can partially avoid the negative effects caused by packet reordering, the additional per-path numbering introduced by LS-SCTP results in an unnecessary overhead, as similar information can be inferred from the sender state variables and SACK chunks in their standard shape. This observation was the underlying principle considered in another loadsharing proposal, called independent per

[9]`http://lksctp.sourceforge.net/`.

path congestion control SCTP (IPCC-SCTP) introduced by Ye et al. [2004]. IPCC-SCTP, instead of using explicit per-path numbering (as LS-SCTP), provides a local per-path mapping for each SCTP packet. This information is only necessary at the sender-side to control the congestion; thus, the sending of extra information is avoided. Thanks to this local mapping, IPCC-SCTP can govern congestion control, SACK processing, and retransmission handling on each path separately instead of doing it for the entire association, as in standard SCTP. In any case, the IPCC-SCTP solution still faces some shortcomings, for example, it does not solve the reverse traffic problem (frequent SACK generation on the out-of-order arrivals at the receiver side) that have been addressed in other proposals.

5.3.2. Concurrent Multipath Transfer. IPCC-SCTP's implicit per-path sequence numbering approach has been followed in the design of another loadsharing scheme, concurrent multipath transfer (CMT), fully described by Iyengar et al. [2006]. The idea of CMT was first introduced in Iyengar et al. [2004a]. However, in contrast to IPCC-SCTP, CMT has been further developed in the following years. To accommodate CMT, Iyengar et al. propose a new sender architecture in which each path has a separate virtual buffer to guarantee path independence. This modification preserves TCP friendliness under the assumption that the bottleneck is not shared by the paths. Such a virtual multibuffer structure guarantees path independence as far as transmission is concerned but has its implications on congestion control, and therefore, several changes to standard SCTP have been proposed.

(1) To handle congestion control per path and not per association and thus limit severe reordering problems, a sender cwnd growth algorithm (cwnd update for CMT - CUC) has been proposed. Thus, SACKs updating the cumulative transmission sequence number ACK point (CumTSN) received in-order per path and out-of-order per association increase the cwnd on that path.
(2) Fast retransmission needs slight modifications, as reordering introduced by a CMT sender can provoke unnecessary spurious fast retransmissions with cwnd implications. Elimination of spurious fast retransmissions is handled by the split fast retransmit (SFR) algorithm that takes into account not only SACK information but also the transmission destination for each transmission sequence number (TSN) when triggering the retransmission to a given path.
(3) The CMT receiver should not send immediate SACKs for packets that arrive out of order, as networks may be vulnerable to the increased ACK traffic. As the SCTP receiver does not distinguish loss from reordering introduced by a CMT sender, an algorithm called delayed ACK for CMT (DAC) is used at the sender to correctly infer losses. On the receiver side, the DAC algorithm extends the SACKs with information about the number of data PDUs received since the last SACK was sent.
(4) An appropriate retransmission policy for handling retransmissions is needed. This topic has been investigated in more detail in Iyengar et al. [2004b], with five retransmission policies proposed. As for the bulk applications considered, the best results were achieved by loss rate-based policies that either sent retransmissions to the path with the highest cwnd value or to the path with the highest ssthresh value.

The CMT extension to standard SCTP discussed so far allows congestion control to properly deal with packet reordering. However, packet reordering still leads to an additional problem that impacts the rate control. This problem is referred to as *receiver buffer blocking*, where the receiver buffer is filled with out-of-order data due to complete or short-term failures and, under limited receiver buffer size, can cause rate control to stall new data transmission. Receiver buffer blocking has already been tackled by Iyengar et al. [2005, 2007], who suggest selecting an adequate retransmission policy to reduce this problem. Natarajan et al. [2006] propose a solution that partially mitigates the receiver buffer blocking problem, called CMT Potentially Failed (CMT-PF). CMT-PF marks the path that has experienced a failure (a single timeout) as potentially failed and stops transmitting data on such a path until a positive heartbeat probe is returned. An extensive evaluation of CMT-PF presented in Natarajan et al. [2009] shows that the proposed solution performs better or similar but never worse than CMT. The CMT-PF proposal has been conceived for lossy network scenarios, although not particularly designed with wireless networks in mind, thus facilitating the idea of applying CMT to also improve transport layer handover. In this regard, using loadsharing in such a context was originally proposed by Goff and Phatak [2004], as already stated in Section 5.2.

The impact of the CMT proposal on the research community has been quite substantial, spurring a considerable amount of articles in the following years evaluating possible applications of CMT. These include applications to multihop wireless scenarios [Aydin and Shen 2009] or use of CMT during a handover process [Budzisz et al. 2009; Huang et al. 2009], as already mentioned in Section 5.2.4.

5.3.3. Source Scheduling. Additionally, loadsharing schemes can be complemented with source scheduling algorithms to manage available paths in a more efficient way and thus minimize packet reordering issues. Casetti et al. [2004] provides an initial idea for load balancing based on a bandwidth-aware source-scheduling extension to SCTP. Casetti et al. suggest sending a pair of heartbeat packets back-to-back in order to estimate the available bandwidth and picking the fastest path to transmit data, as the simplest approach named SBPP-SCTP. This idea is further developed with a design of Westwood-like SCTP (W-SCTP) proposed by Fiore et al. [2007] and Fiore and Casetti [2005]. Apart from introducing a multibuffer structure and per-path congestion control—modifications similar to those described earlier for CMT—Fiore et al. employ a packet scheduler that maximizes the chance that packets sent on paths with different bandwidths will arrive in order at the receiver, thus minimizing the receiver buffer blocking problem. The bandwidth estimation is made in a Westwood-like manner, giving the name for the proposed scheme. Moreover, an explicit advance acknowledgment algorithm (to avoid a problem similar to that of HoL blocking in which packets are chosen to be transmitted on an inactive path) and a minimum bandwidth estimate threshold (to avoid unused paths becoming unavailable) are provided to increase the robustness of the presented approach. An exhaustive comparison of both approaches (SBPP and W-SCTP) based on emulation results is provided by Perotto et al. [2007]. The overall conclusion is that the Westwood-like estimate works better than the pair of packets (SBPP), especially with interfering traffic, as it is able to estimate current available bandwidth rather than just the link capacity. W-SCTP is also used together with the PR-SCTP extension (the entire scheme is named W-SCTP-PR) to provide support for real-time applications, as presented in Fiore and Casetti [2005]. Real-time traffic poses additional constraints on low-delay jitter values and out-of-sequence packets that are shown to be met by the W-SCTP-PR scheme. Further studies on multimedia traffic have also been conducted by Rossi et al. [2006].

5.3.4. Classification of Related Articles. Figure 9 presents the distribution of loadsharing-related research over the remaining dimensions of the taxonomy. Again, the most common application is bulk transfer. However, contributions are found in each category except in Other Applications. There is also considerable versatility in the analyzed network scenarios, more or less equally distributed between the wired and wireless domain. Simulation proves again to be the most common study approach, whereas

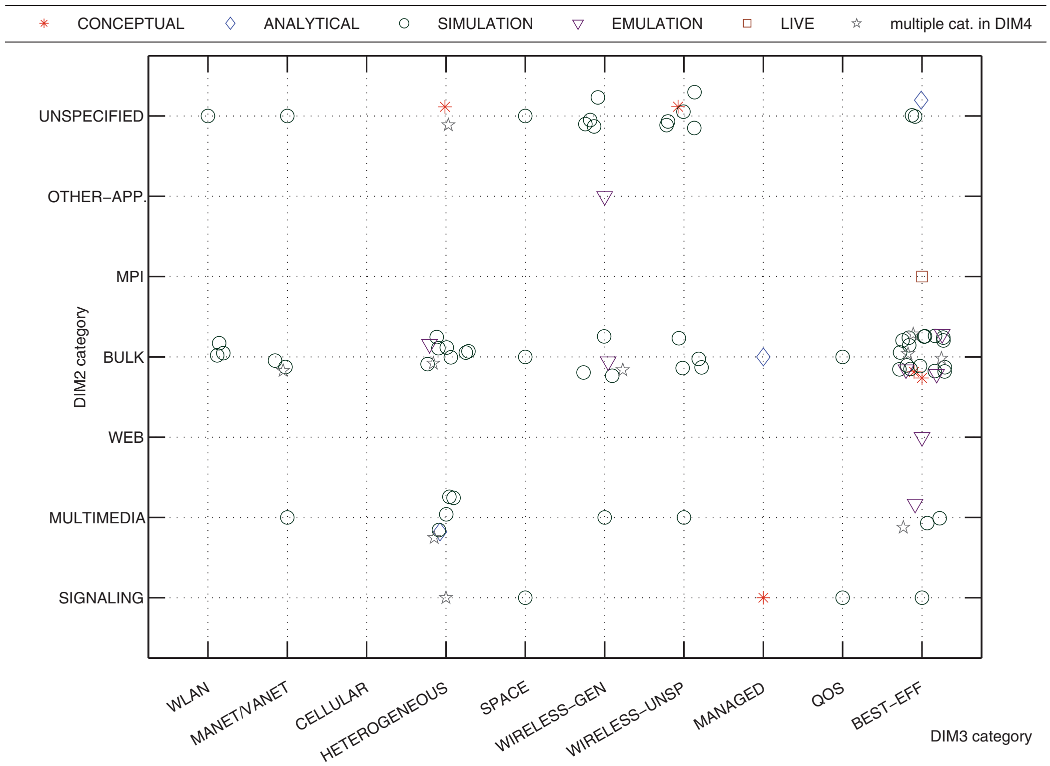

Fig. 9. Scatter plot of all loadsharing-related articles.

articles based on emulation results are the second main group but with a considerably lower share.

6. CONCLUSION

The SCTP protocol, which was initially designed for transporting signaling messages over IP networks, has had a notable impact on the research community during the first ten years since its standardization in 2000. SCTP is now an established general transport protocol, and the presented taxonomy and classification of more than 430 SCTP-related articles clearly show that SCTP has considerable potential for application in diverse fields. Moreover, the number of Internet drafts, protocol implementations for major OSes, and simulator models illustrate the dynamics of the SCTP research community. From a deployment perspective, as a new transport protocol, SCTP has to challenge the well-established transport protocols in the TCP/IP protocol stack, which seems extremely difficult without having a strong motivating application. In this regard, SCTP is progressively being adopted as a key protocol stack component for signaling transfer within VoIP architectures and next-generation mobile networks. However, out of such a signaling transport context, SCTP has not yet been able to find its killer application that would drive its adoption. To identify possible fields where such an application could be developed, this article provided a comprehensive overview of SCTP research published so far. The proposed taxonomy has proved a convenient means of grouping the available research and giving insight on the distribution of SCTP research between the different protocol features, application areas, networking environments, and research approaches. As a key observation, the conducted analysis has revealed that SCTP research centers around two new protocol features: multihoming and multistreaming—addressed by about 56% and 18% of all the analyzed articles, respectively. Attending to the relevance of the multihoming feature, a detailed survey explored in greater depth the published research works related to the exploitation of the multihoming feature, using our proposed taxonomy as a basis for the discussion. In particular, focus was placed on the initial use for multihoming robustness, as well as on the later use for transport layer handover and loadsharing.

REFERENCES

Abd El Al, A., Saadawi, T., and Lee, M. 2004a. Improving throughput and reliability in mobile wireless networks via transport layer bandwidth aggregation. *Comput. Netw. 46,* 5, 635–649.

Abd El Al, A., Saadawi, T., and Lee, M. 2004b. LS-SCTP: A bandwidth aggregation technique for stream control transmission protocol. *Comput. Commun. 27,* 10, 1012–1024.

Argyriou, A. and Madisetti, V. 2007. The design and evaluation of an end-to-end handoff management protocol. *Wirel. Netw. 13,* 1, 61–75.

Aydin, I., Seok, W., and Shen, C.-C. 2003. Cellular SCTP: A transport-layer approach to internet mobility. In *Proceedings of the 12th International Conference on Computer Communications and Networks (ICCCN'03)*. 285–290.

Aydin, I. and Shen, C.-C. 2005. Evaluating cellular SCTP over one-hope wireless networks. In *Proceedings of the 62nd IEEE Vehicular Technology Conference (VTC'05)*. Vol. 2. 826–830.

Aydin, I. and Shen, C.-C. 2009. Performance evaluation of concurrent multipath transfer using SCTP multihoming in multihop wireless networks. In *Proceedings of the 8th IEEE International Symposium on Network Computing and Applications (NCA'09)*, 234–241.

Barile, I. 2004. Stream-control transmission protocol. *Dr. Dobb's J. 29,* 9, 32–36.

Bokor, L., Huszak, A., and Jeney, G. 2009. On SCTP multihoming performance in native IPv6 UMTS-WLAN environments. In *Proceedings of the 5th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom'09)*.

Budzisz, Ł., Ferrús, R., Brunstrom, A., Grinnemo, K.-J., Fracchia, R., Galante, G., and Casadevall, F. 2008. Towards transport-layer mobility: Evolution of SCTP multihoming. *Comput. Commun. 31,* 5, 980–998.

Budzisz, Ł., Ferrús, R., and Casadevall, F. 2006. SCTP multihoming performance in dynamically changing channels with the influence of link-layer retransmissions. In *Proceedings of the 64th IEEE Vehicular Technology Conference (VTC'06)*. 2624–2628.

Budzisz, Ł., Ferrús, R., and Casadevall, F. 2009a. Design principles and performance evaluation of mSCTP-CMT for transport-layer based handover. In *Proceedings of the IEEE Vehicular Technology Conference (VTC'09)*.

Budzisz, Ł., Ferrús, R., Casadevall, F., and Amer, P. 2009b. On concurrent multipath transfer in SCTP-based handover scenarios. In *Proceedings of the IEEE International Conference on Communications (ICC'09)*.

Caro, Jr., A., Amer, P., and Stewart, R. 2003a. Retransmission policies with transport layer multihoming. In *Proceedings of the 11th IEEE International Conference on Networks (ICON'03)*. 255–260.

Caro, Jr., A., Amer, P., and Stewart, R. 2003b. Transport layer multihoming for fault tolerance in FCS networks. In *Proceedings of the IEEE Military Communications Conference (MilCom'03)*. 949–953.

Caro, Jr., A., Amer, P., and Stewart, R. 2004a. End-to-end failover thresholds for transport layer multihoming. In *Proceedings of the IEEE Military Communications Conference (MilCom'04)*. 99–105.

Caro, Jr., A., Amer, P., and Stewart, R. 2004b. Retransmission schemes for end-to-end failover with transport layer multihoming. In *Proceedings of the IEEE Global Telecommunications Conference (GlobeCom'04)*. Vol. 3. 1341–1347.

Caro, Jr., A., Amer, P., and Stewart, R. 2006a. Rethinking end-to-end failover with transport layer multihoming. *Ann. Telecommun. 61,* 1-2, 92–114.

Caro, Jr., A., Amer, P., and Stewart, R. 2006b. Retransmission policies for multihomed transport protocols. *Comput. Commun. 29,* 10, 1798–1810.

Casetti, C., Chiasserini, C. F., Fracchia, R., and Meo, M. 2006. AISLE: Autonomic interface selection for wireless users. In *Proceedings of the International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM'06)*. 42–48.

Casetti, C., Chiasserini, C. F., Fracchia, R., and Meo, M. 2008. Autonomic interface selection for mobile wireless users. *IEEE Trans. Veh. Technol. 57,* 6, 3666–3678.

CASETTI, C., GALANTE, G., AND GRECO, R. 2004. Load balancing over multipaths using bandwidth-aware source scheduling. In *Proceedings of the 7th International Symposium on Wireless Personal Multimedia Communications (WPMC'04)*.

CHANG, L.-H., LIN, H.-J., AND CHANG, I. 2007. Dynamic handover mechanism using mobile SCTP in contention based wireless network. In *Proceedings of the 5th International Symposium on Parallel and Distributed Processing and Applications (ISPA'07)*. Lecture Notes in Computer Science, vol. 4742, Springer, Berlin, 821–831.

CHANG, M., LEE, M., AND KOH, S. J. 2004. Transport layer mobility support utilizing link signal strength information. *IEICE Trans. Commun. E87-B,* 9, 2548–2556.

EKLUND, J., BRUNSTROM, A., AND GRINNEMO, K.-J. 2008. On the relation between SACK delay and SCTP failover performance for different traffic distributions. In *Proceedings of the 5th International Conference on Broadband Communications, Networks, and Systems (BROADNETS'08)*. 577–584.

EKLUND, J., GRINNEMO, K.-J., BAUCKE, S., AND BRUNSTROM, A. 2010. Tuning SCTP failover for carrier grade telephony signaling. *Comput. Netw. 54,* 1, 133–149.

FIORE, M. AND CASETTI, C. 2005. An adaptive transport protocol for balanced multihoming of real-time traffic. In *Proceedings of the IEEE Global Telecommunications Conference (GlobeCom'05)*. Vol. 2. 1091–1096.

FIORE, M., CASETTI, C., AND GALANTE, G. 2007. Concurrent multipath communication for real-time traffic. *Comput. Commun. 30,* 17, 3307–3320.

FITZPATRICK, J., MURPHY, S., ATIQUZZAMAN, M., AND MURPHY, J. 2009. Using cross-layer metrics to improve the performance of end-to-end handover mechanisms. *Comput. Commun. 32,* 15, 1600–1612.

FITZPATRICK, J., MURPHY, S., AND MURPHY, J. 2006. SCTP based handover mechanism for VoIP over IEEE 802.11b wireless LAN with heterogeneous transmission rates. In *Proceedings of the IEEE International Conference on Communications (ICC'06)*. Vol. 5. 2054–2059.

FRACCHIA, R., CASETTI, C., CHIASSERINI, C.-F., AND MEO, M. 2005. A WiSE extension of SCTP for wireless networks. In *Proceedings of the IEEE International Conference on Communications, (ICC'05)*. 1448–1453.

FRACCHIA, R., CASETTI, C., CHIASSERINI, C.-F., AND MEO, M. 2007. WiSE: Best-path selection in wireless multihoming environments. *IEEE Trans. Mob. Comput. 6,* 10, 1130–1141.

FU, S. AND ATIQUZZAMAN, M. 2004. SCTP: State of the art in research, products, and technical challenges. *IEEE Commun. Mag. 42,* 4, 64–76.

GARCIA, J. AND BUDZISZ, Ł. 2010. SCTP Database webpage. `http://www.kau.se/sctpsurvey`.

GAUCH, A. AND NISHIDA, Y. 2006. SCTP profiling framework for multi-homed environment. In *Proceedings of the 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking (MobiShare'06)*. 54–56.

GOFF, T. AND PHATAK, D. 2004. Unified transport layer support for data striping and host mobility. *IEEE J. Selec. Areas Commun. 22,* 4, 737–746.

GRINNEMO, K.-J. AND BRUNSTROM, A. 2004. Performance of SCTP-controlled failovers in M3UA-based SIGTRAN networks. In *Proceedings of the Advanced Simulation Technologies Conference (ASTC'04)*.

GRINNEMO, K.-J. AND BRUNSTROM, A. 2005. Impact of traffic load on SCTP failovers in SIGTRAN. In *Proceedings of the International Conference on Networking 2005 (ICN'05)*. 774–783.

HAYES, D. A., BUT, J., AND ARMITAGE, G. 2009. Issues with network address translation for SCTP. *SIGCOMM Comput. Commun. Rev. 39,* 1, 23–33.

HEINZ, G. AND AMER, P. 2004. Priorities in stream transmission control protocol (SCTP) multistreaming. In *Proceedings of the 8th World Multiconference on Systemics, Cybernetics, and Informatics (SCI'04)*.

HONDA, M., NISHIDA, Y., NAKAZAWA, J., AND TOKUDA, H. 2007. Performance enhancement of transport layer handover on single-homed mobile nodes. *IEICE Trans. Commun. E90-B,* 10, 2683–2692.

HUANG, C.-M., LIN, M.-S., CHANG, L.-H., AND CHANG, W.-Y. 2009. An MIH-assisted handoff mechanism for concurrent multipath transfer in wireless multihomed networks. In *Proceedings of the IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*.

HUANG, C. M. AND TSAI, C. H. 2007. WiMP-SCTP: Multi-path transmission using stream control transmission protocol (SCTP) in wireless networks. In *Proceedings of the 21st International Conference on Advanced Information Networking and Applications Workshops/Symposia (AINAW'07)*. Vol. 2. 209–214.

IYENGAR, J., AMER, P., AND STEWART, R. 2004a. Concurrent multipath transfer using transport layer multihoming—performance under varying bandwidth proportions. In *Proceedings of the IEEE Military Communications Conference (MilCom'04)*. 238–244.

IYENGAR, J., AMER, P., AND STEWART, R. 2004b. Retransmission policies for concurrent multipath transfer using SCTP multihoming. In *Proceedings of the 12th IEEE International Conference on Networks (ICON'04)*. 713–719.

IYENGAR, J., AMER, P., AND STEWART, R. 2005. Receive buffer blocking in concurrent multipath transfer. In *Proceedings of the IEEE Global Telecommunications Conference (GlobeCom'05)*. Vol. 1. 121–126.

IYENGAR, J., AMER, P., AND STEWART, R. 2006. Concurrent multipath transfer using SCTP multihoming over independent end-to-end paths. *IEEE/ACM Trans. Netw. 14,* 5, 951–964.

IYENGAR, J., AMER, P., AND STEWART, R. 2007. Performance implications of a bounded receive buffer in concurrent multipath transfer. *Comput. Commun. 30,* 4, 818–829.

JUNGMAIER, A. AND RATHGEB, E. 2006. On SCTP multi-homing performance. *Telecommun. Syst. 31,* 2-3, 141–161.

JUNGMAIER, A., RATHGEB, E., SCHOPP, M., AND TUEXEN, M. 2001. SCTP—a multi-link end-to-end protocol for IP-based networks. *AEU Int. J. Electron. Commun. 55,* 1, 46–54.

JUNGMAIER, A., RATHGEB, E., AND TUEXEN, M. 2002. On the use of SCTP in failover scenarios. In *Proceedings of the 6th World Multiconference on Systemics, Cybernetics, and Informatics (SCI'02)*. 363–368.

KASHIHARA, S., IIDA, K., KOGA, H., KADOBAYASHI, Y., AND YAMAGUCHI, S. 2004. Multi-path transmission algorithm for end-to-end seamless handover across heterogeneous wireless access networks. *IEICE Trans. Commun. E87-B,* 3, 490–496.

KIM, D. P., KOH, S. J., AND KIM, S. W. 2006. mSCTP-DAC: Dynamic address configuration for msctp handover. In *Proceedings of the International Conference on Embedded and Ubiquitous Computing (EUC'06)*. Lecture Notes in Computer Science, vol. 4096, Springer, Berlin, 244–253.

KIM, S. T., KOH, S. J., AND KIM, Y. J. 2007. Performance of SCTP for IPTV applications. In *Proceedings of the 9th International Conference on Advanced Communication Technology (ICACT'07)*. Vol. 3. 2176–2180.

KOH, S. J., CHANG, M. J., AND LEE, M. 2004. mSCTP for soft handover in transport layer. *IEEE Commun. Lett. 8,* 3, 189–191.

KRISHNAN, S., MONTAVONT, N., NJEDJOU, E., VEEREPALLI, S., AND YEGIN, A. 2007. RFC 4957, Link-layer event notifications for detecting network attachments. `http://www.ietf.org/rfc/rfc4957.txt`.

MA, L., YU, F., AND LEUNG, V. 2007. Performance improvements of mobile SCTP in integrated heterogeneous wireless networks. *IEEE Trans. Wirel. Commun. 6,* 10, 3567–3577.

MA, L., YU, F., LEUNG, V., AND RANDHAWA, T. 2004. A new method to support UMTS/WLAN vertical handover using SCTP. *IEEE Wirel. Commun. 11,* 4, 44–51.

NATARAJAN, P., EKIZ, N., AMER, P., AND STEWART, R. 2009. Concurrent multipath transfer during path failure. *Comput. Commun. 32,* 15, 1577–1587.

NATARAJAN, P., IYENGAR, J., AMER, P., AND STEWART, R. 2006. Concurrent multipath transfer using transport layer multihoming: Performance under network failures. In *Proceedings of the IEEE Military Communications Conference (MilCom'06)*.

NOONAN, J., PERRY, P., AND MURPHY, J. 2004. Client controlled network selection. In *Proceedings of the 5th IEE International Conference on 3G Mobile Communication Technologies*. 543–547.

NOONAN, J., PERRY, P., AND MURPHY, J. 2006. End-point synchronisation and handover for multi-homed services. *IEE Proc. Commun. 153,* 5, 691–696.

NOONAN, J., PERRY, P., MURPHY, S., AND MURPHY, J. 2006. Stall and path monitoring issues in SCTP. In *Proceedings of the 25th IEEE Conference on Computer Communications (InfoCom'06)*.

PAXON, V., ALLMAN, M., AND STEVENS, W. 1999. RFC 2581, TCPs congestion control. `http://www.ietf.org/rfc/rfc2581.txt`.

PEROTTO, F., CASETTI, C., AND GALANTE, G. 2007. SCTP-based transport protocols for concurrent multipath transfer. In *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC'07)*. 2971–2976.

QIAO, Y., FALLON, E., MURPHY, L., MURPHY, J., HANLEY, A., ZHU, X., MATTHEWS, A., CONWAY, E., AND HAYES, G. 2007. SCTP performance issue on path delay differential. In *Proceedings of the 5th International Conference on Wired/Wireless Internet Communications (WWIC'07)*. Lecture Notes in Computer Science, vol. 4517, Springer, Berlin, 43–54.

RIEGEL, M. AND TUEXEN, M. 2007. Mobile SCTP, IETF draft. `http://tools.ietf.org/id/draft-riegel-tuexen-mobile-sctp-09.txt`. Work in progress.

ROSSI, C., CASETTI, C., FIORE, M., AND SCHONFELD, D. 2006. A partially reliable transport protocol for multiple-description real-time multimedia traffic. In *Proceedings of the IEEE International Conference on Image Processing*. 1301–1304.

SAMTANI, S., IYENGAR, J., AND FECKO, M. 2003. SCTP multistreaming: preferential treatment among streams. In *Proceedings of the IEEE Military Communications Conference (MilCom'03)*. 966–970.

SHEU, T. L. AND TU, C. M. 2007. Dynamic chunk adjustment for prioritized SCTP video streaming in heterogeneous WLAN and UMTS networks. In *Proceedings of the International Wireless Communications and Mobile Computing Conference (IWCMC'07)*. 612–617.

STEWART, R. 2007. RFC 4960, Stream Control Transmission Protocol (SCTP). `http://www.ietf.org/rfc/rfc4960.txt`.

STEWART, R. AND AMER, P. 2007. Why is SCTP needed given TCP and UDP are widely available? `http://www.isoc.org/briefings/017/`.

STEWART, R., ARIAS-RODRIGUEZ, I., POON, K., CARO, JR., A., AND TUEXEN, M. 2006. RFC 4460, Stream Control Transmission Protocol (SCTP) specification errata and issues. `http://www.ietf.org/rfc/rfc4460.txt`.

STEWART, R. AND METZ, C. 2001. SCTP—New transport protocol for TCP/IP. *IEEE Internet Comput. 5,* 6, 64–69.

STEWART, R., RAMALHO, M., XIE, Q., TUEXEN, M., AND CONRAD, P. 2004. RFC 3758, Stream Control Transmission Protocol (SCTP) Partial Reliability Extension. `http://www.ietf.org/rfc/rfc3758.txt`.

STEWART, R., TUEXEN, M., AND CAMARILLO, G. 2007. RFC 5062, Security attacks found against the Stream Control Transmission Protocol (SCTP) and current countermeasures. `http://www.ietf.org/rfc/rfc5062.txt`.

STEWART, R., XIE, Q., MORNEAULT, K., SHARP, C., SCHWARZBAUER, H., TAYLOR, T., RYTINA, I., KALLA, M., ZHANG, L., AND PAXSON, V. 2000. RFC 2960, Stream Control Transmission Protocol (SCTP). `http://www.ietf.org/rfc/rfc2960.txt`. Obsoleted by RFC 4960.

STEWART, R., XIE, Q., TUEXEN, M., MARUYAMA, S., AND KOZUKA, M. 2007. RFC 5061, Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration. `http://www.ietf.org/rfc/rfc5061.txt`.

STONE, J., STEWART, R., AND OTIS, D. 2002. RFC 3309, Stream Control Transmission Protocol (SCTP) checksum change. `http://www.ietf.org/rfc/rfc3309.txt`. Obsoleted by RFC 4960.

TUEXEN, M., RUNGELER, I., STEWART, R., AND RATHGEB, E. 2008. Network address translation for the stream control transmission protocol. *IEEE Netw. 22,* 5, 26–32.

TUEXEN, M., STEWART, R., LEI, P., AND RESCORLA, E. 2007. RFC 4895, Authenticated chunks for the Stream Control Transmission Protocol (SCTP). `http://www.ietf.org/rfc/rfc4895.txt`.

YE, G., SAADAWI, T., AND LEE, M. 2004. IPCC-SCTP: An enhancement to the standard SCTP to support multihoming efficiently. In *Proceedings of the 23rd IEEE International Conference on Performance, Computing, and Communications (IPCCC'04)*. 523–530.

YOO, S. K., KIM, H. G., AND SOHN, S. 2002. Enhancement of failover using application layer watchdog and SCTP heartbeat in diameter. In *Proceedings of the CDMA International Conference (CIC'02)*. Lecture Notes in Computer Science, vol. 2524, Springer, Berlin, 239–246.

ZEADALLY, S. AND SIDDIQUI, F. 2007. An empirical analysis of handoff performance for SIP, mobile IP, and SCTP protocols. *Wirel. Pers. Commun. 43,* 2, 589–603.

Received February 2010; revised October 2010; accepted January 2011