SIP Conferencing

Henning Schulzrinne
Dept. of Computer Science
Columbia University
New York, New York
(sip:)schulzrinne@cs.columbia.edu

Conference International SIP – Paris, France

February 21, 2001

Joint work with Gautam Nair, Jonathan Rosenberg and Kundan Singh

Overview

- conferencing models
- full-mesh conferencing
- implementation of a conferencing server
- open issues

Conference models: end system mixing

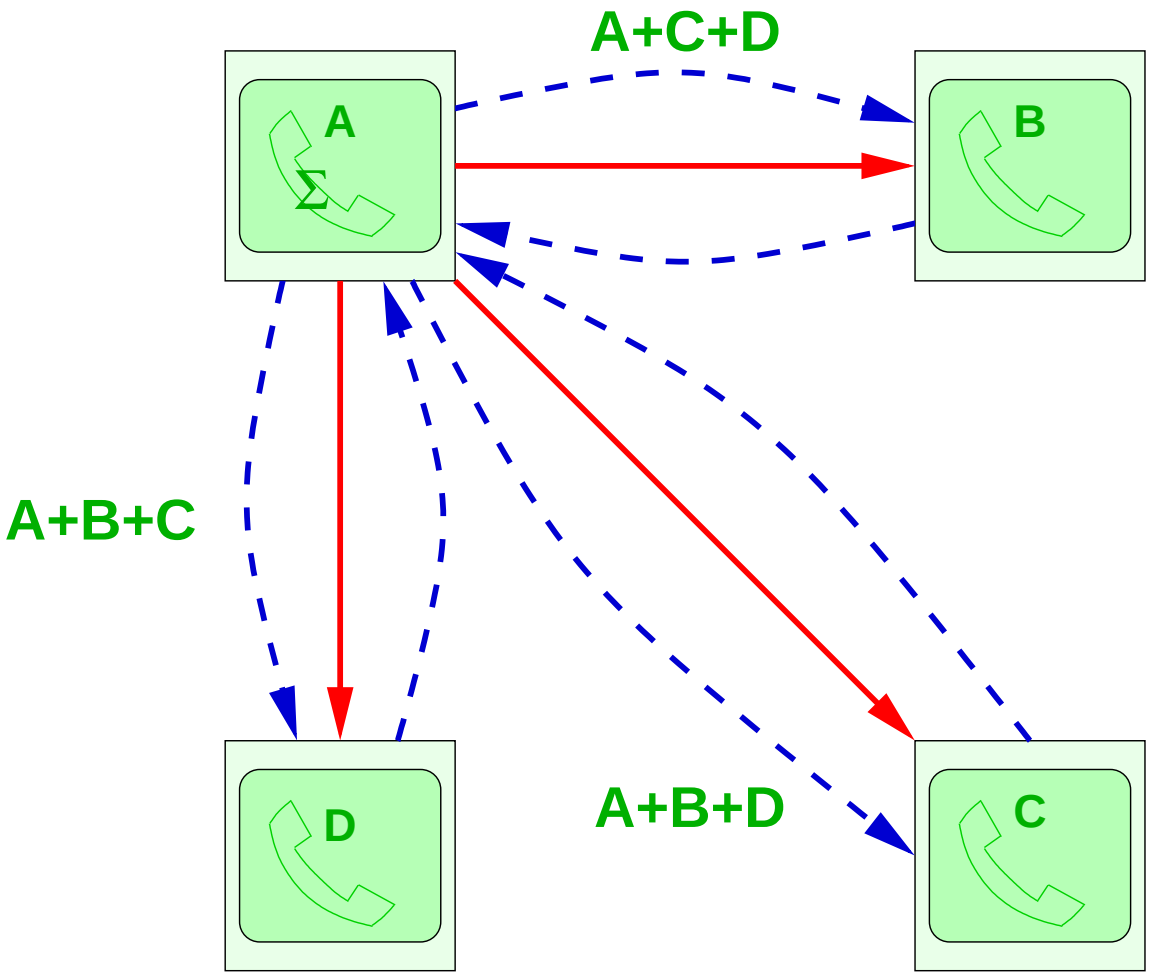

Conference models: multicast transmit & receive

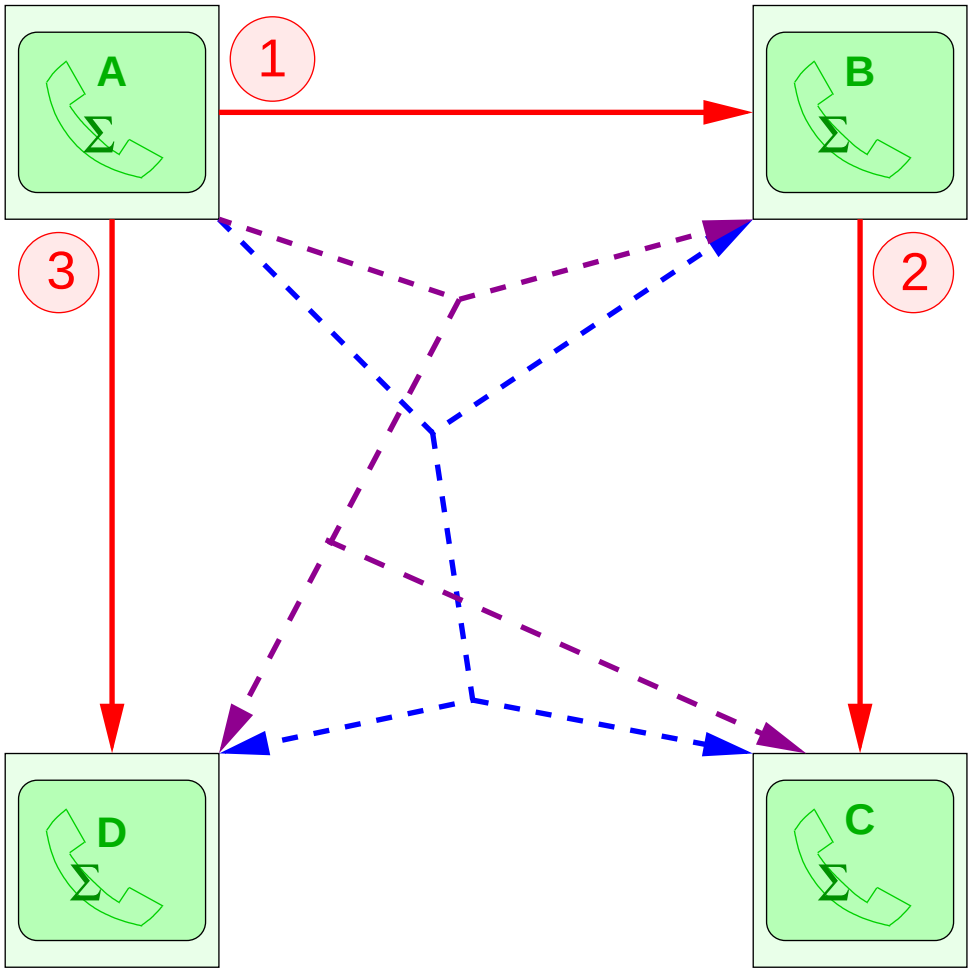

May be suitable for small-scale multicast proposals (“xcast”)

Conference models: multicast receive, unicast transmit

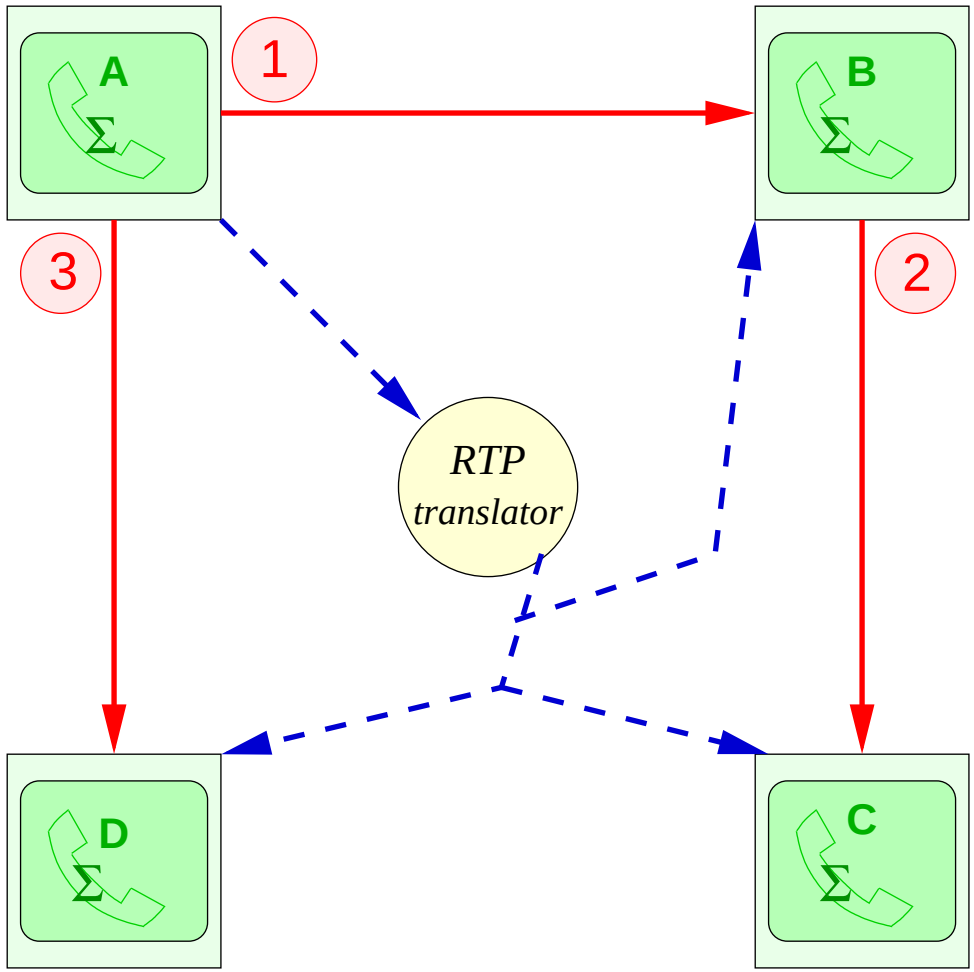

Conference models: central server

can be call-out or dial-in

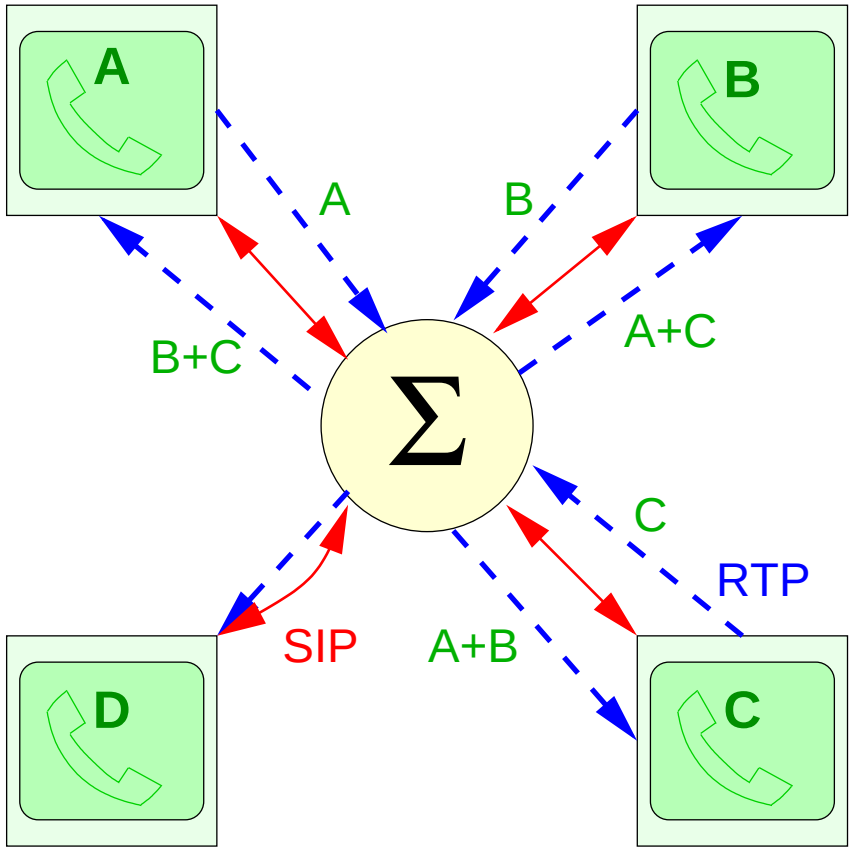

Conference models: central server

- conference identified by SIP URL, e.g., sip2001@upperside.fr
- ad-hoc conferences useful for three-party calls
- can create ad-hoc conference by using random URL, e.g., sip:a7hytaskp09878a@a.servers.com, use REFER to get others to add themselves to that conference
- for ad-hoc, conference lasts until last one leaves

Conference models: decentralized

Central signaling server, mesh media

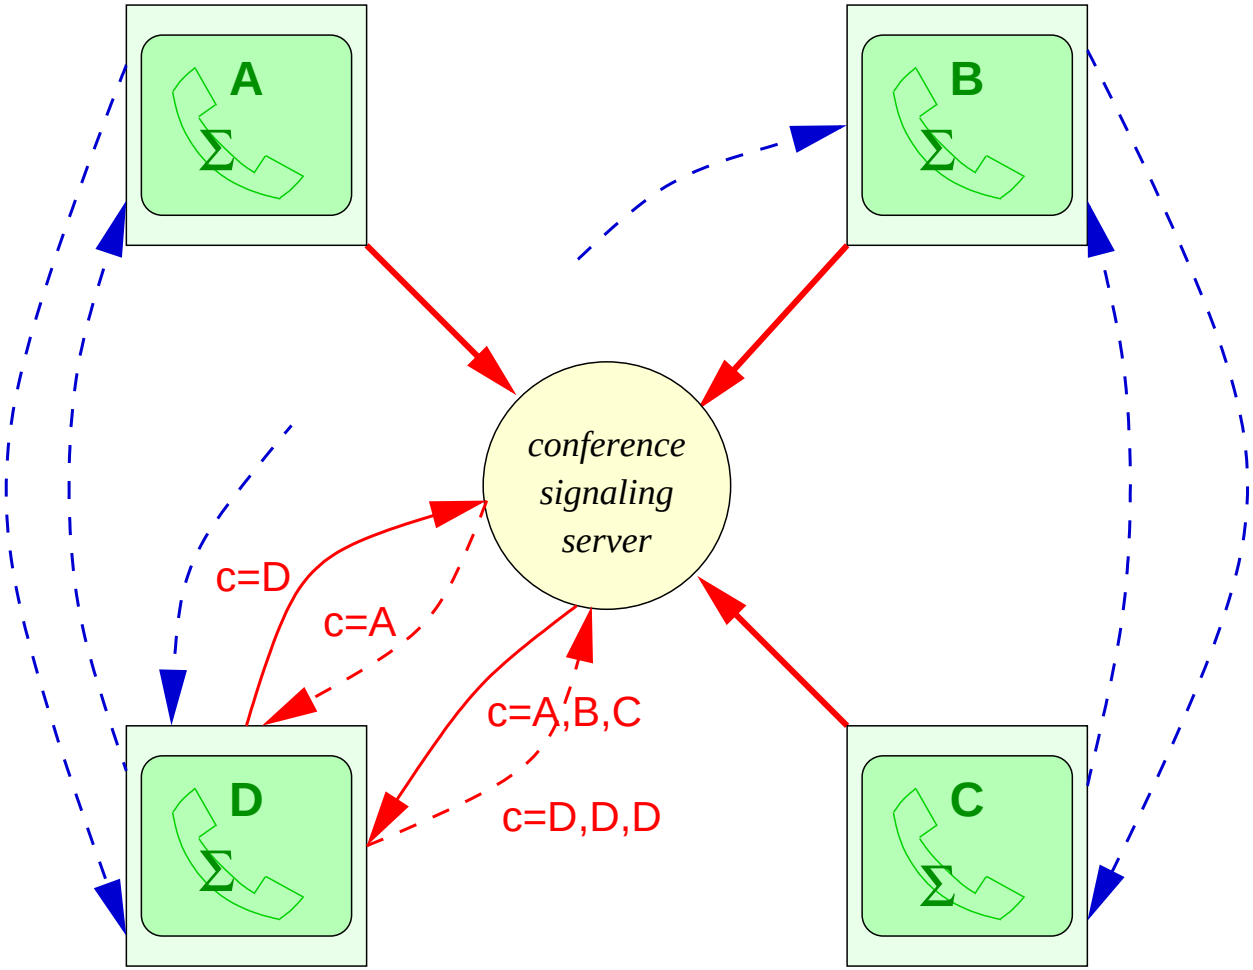

Conference models: full mesh

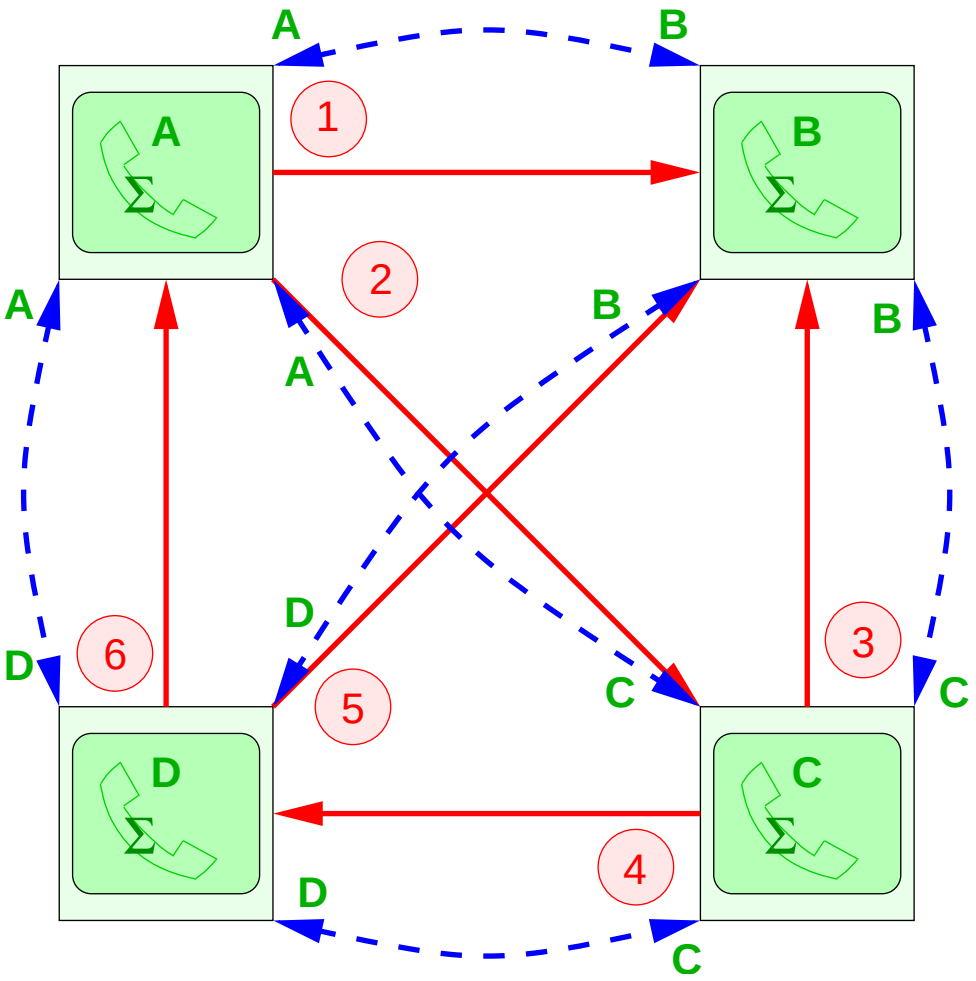

A invites B
A invites C
C invites B
C invites D
D invites A, B

Full mesh

- +: all participants are equal
- +: easy to control media distribution
- need to deal with simultaneous call-outs of two parties
- departures while other party joins
- UAC and UAS send membership list in request and 200
- if either learns about a new member, send INVITE
- allows joining of separated clouds
- departing member sends BYE to everyone – may need to refuse INVITEs from new arrivals (partitions!)

Conference models: full mesh

A invites C and B invites D *simultaneously*:

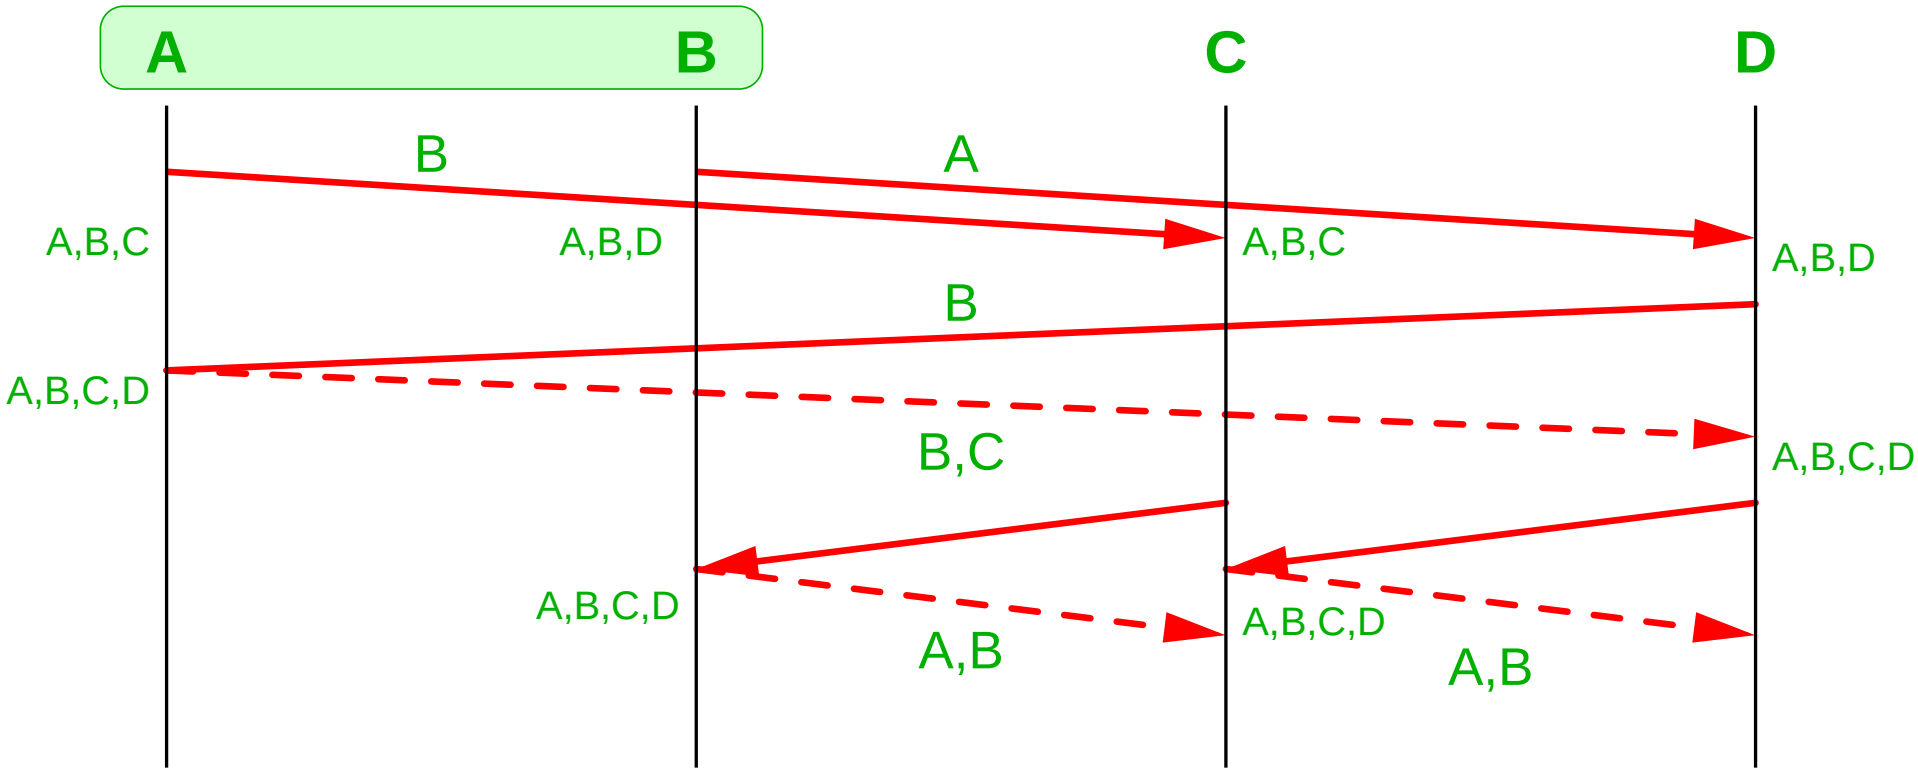

Full mesh: open issues

- conference identification?
- how to identify that call legs are related (should be mixed)
- call-in for mesh?
- refresh mechanism to heal network partitions

Conference models – a comparison

Name	SIP	media	adding	joining	members	scale
End-Mixing	tree	tree	normal INVITE	normal INVITE	RTCP	small
Multicast	pairs	mcast	normal INVITE	mcast join	RTCP	large
Dial-In	star	star	REFER	normal INVITE	RTCP	medium
Ad-Hoc	star	star	REFER	normal INVITE	RTCP	medium
Dial-Out	star	star	REFER	normal INVITE	RTCP	medium
Decentral	star	mesh	refer +	normal	RTCP	medium
Mesh	mesh	mesh	INVITE	?	SIP	medium

Conference models – complexity

M active senders, N participants

Properties	central	full mesh	mcast	uni rx, mcast tx	end mixing
Topology	Star	full mesh	m-cast tree	star + mcast tree	ad-hoc
Server proc.	$O(M+N)$	n/a	n/a	$O(M+N)$	n/a
Endpoint proc.	$O(1)$	$O(M)$	$O(M)$	$O(1)$	variable
Server bw	$O(M+N)$	n/a	n/a	$O(M)$	n/a
Endpoint bw	$O(1)$	$O(M)$	$O(1)$	$O(1)$	variable
Scaling	medium	medium	large	large	medium
Heterogen. UA	yes	yes	no	no	yes (partially)
Own media?	no	no	no	yes	no

Columbia sipconf conferencing server

- central-server model for mixing
- commodity hardware, easily scalable to rack full (conferences: avg. 5, 45 minute duration)
- mixes audio streams, replicates RTP (and UDP) streams
- works for video and text chat (RFC 2793)

Mixing heterogeneous streams

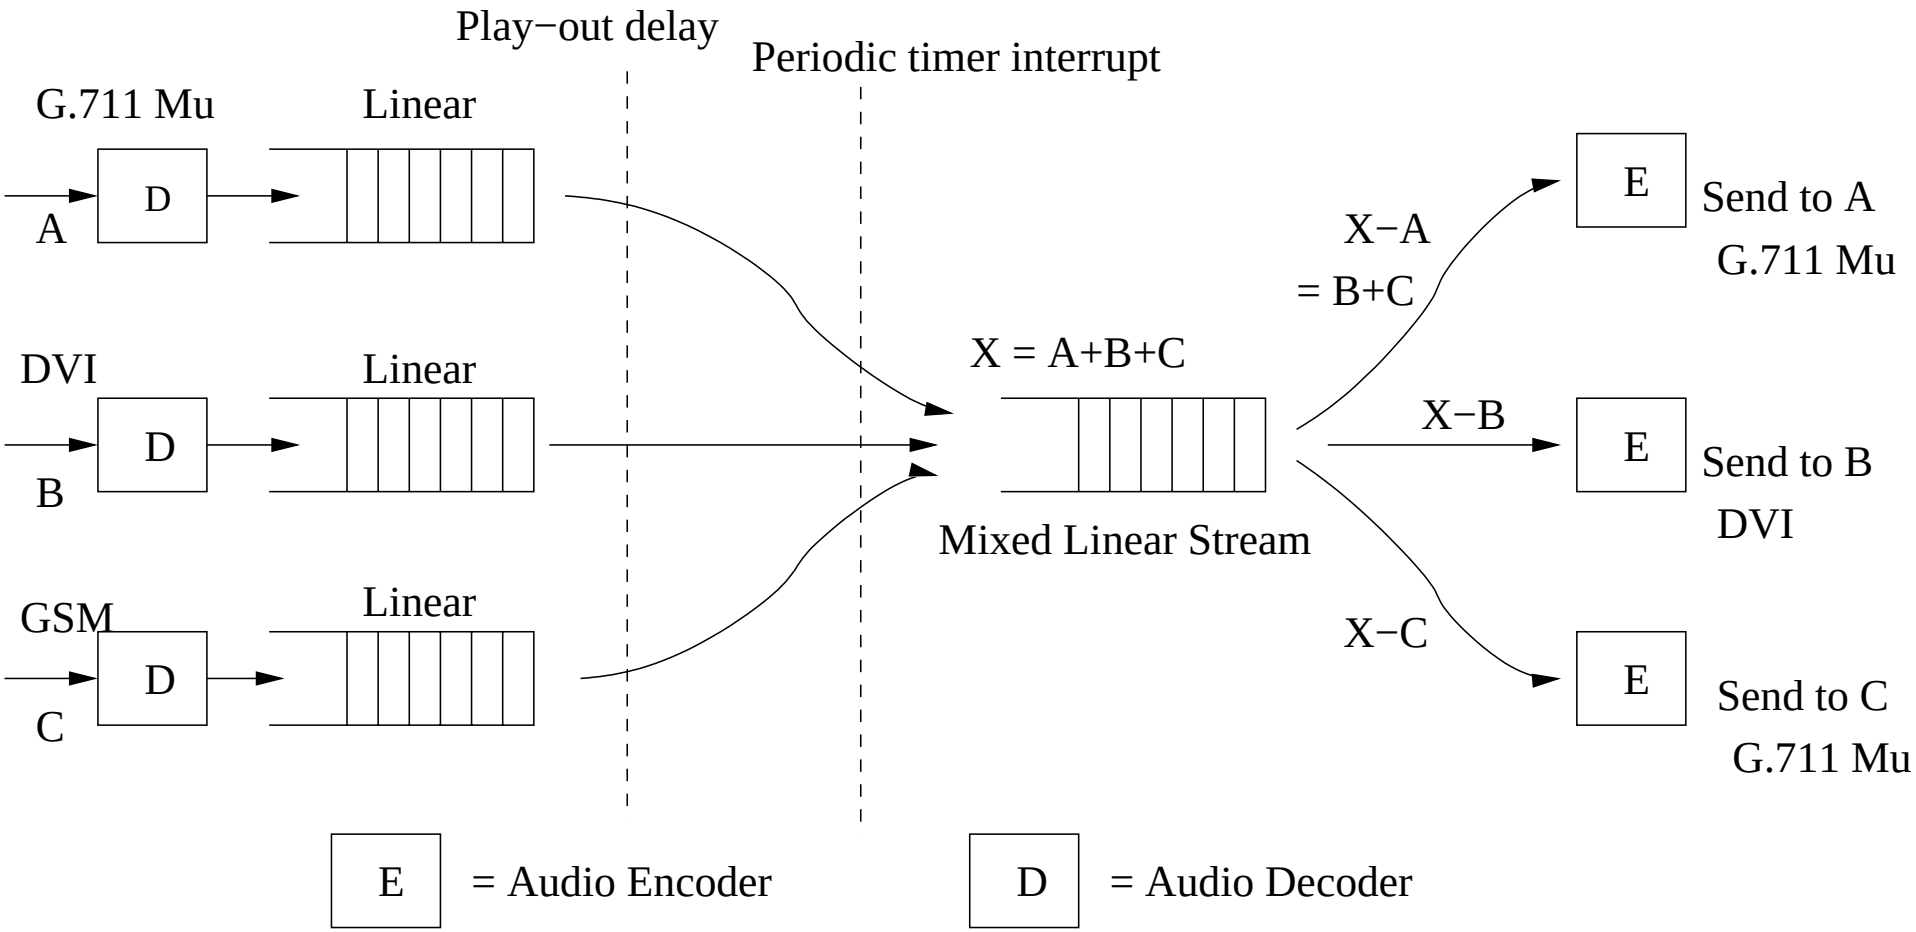

Example: Columbia software conference server

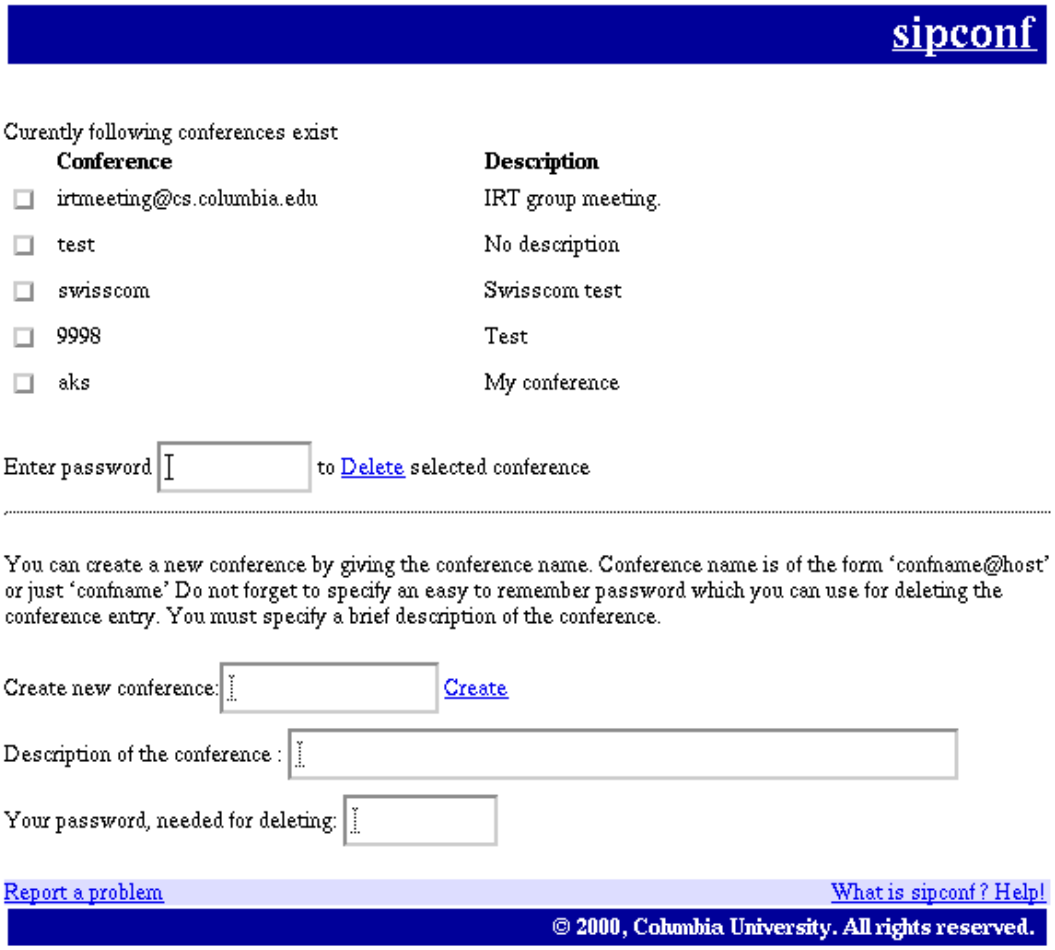

Performance for single conference

SPARC Ultra 10 with 350 MHz CPU:

Participants	CPU (%)	memory (MB)	bandwidth (Mb/s) inbound	bandwidth (Mb/s) outbound
2	< 0.1	2.7	0.08	0.07
20	< 1	6.0	0.08	1.37
40	2-3	9.6	0.08	2.81
60	5	13	0.08	4.25
80	10-15	17	0.08	5.69
100	35-50	22	0.08	7.13
120	50-70	26	0.08	8.59

Performance for three-party conferences

Good quality up to 15 conferences

Confer-ences	partici-pants	CPU (%)	memory (MB)	bandwidth (Mb/s) inbound	outbound
3	9	< 0.4	4.1	0.72	0.65
6	18	< 2.0	5.7	1.44	1.30
9	27	7-13	7.3	2.16	1.94
12	36	15-20	9	2.88	2.60
15	45	25	10	3.60	3.24
18	54	30	12	4.32	3.89

Conclusion

- need many different models of conferencing
- trade-off infrastructure vs. complexity vs. scaling
- mostly, can be handled by existing SIP mechanisms
- full mesh requires further work
- transition between conferencing models?