

MP-DCCP Interoperability Test with DT

Mingxuan Teng, Minjun Xi
from xiaomi

What we did

We ported MP-DCCP to Android 13 (Android T) phones with 5.10 Kernel, and verified the availability and interoperability of MP-DCCP through LAN and WAN iperf.

- Read MP-DCCP specification
- Ported MP-DCCP & U-DCCP converter (see Backup)
- Solved compilation problems (see Backup)
- iperf interoperability test

Preparation of equipment and environment

The interoperability test is based on the MP-DCCP Linux reference implementation , using [iperf3 with MP-DCCP support](#), both LAN and WAN environments are setup for test.

- Test equipment

Mobile phone: Redmi K50G with Dual-WLAN capability, running **Xiaomi modified** [5.10.101-version](#)

Virtual machine: Kernel version is [4.14.111](#)

Remote server: Kernel version is [4.14.111](#)

- Test environment

1. LAN

Virtual machine (server) IP: 192.168.3.xx

Mobile phone (client) IP: 192.168.3.xx & 192.168.31.xx

2. WAN

Remote server (server) IP:213.239.223.XXX

Mobile phone (client) IP:1.202.162.XXX

UDP <-> DCCP header conversion used to avoid problems with middle boxes



<https://datatracker.ietf.org/doc/html/draft-amend-tsvwg-dccp-udp-header-conversion-01>

Preparation for Interoperability test

Configure the MP-DCCP-related routing rules and interfaces, and specify one port known by both parties to convert the UDP protocol to the MP-DCCP to ensure the normal transmission of data packets

- Options that need to be configured

1. Enable MP-DCCP on every physical interface
2. Configure transfer options
3. Configure routing on every physical interface
4. Set U-DCCP translation port

- Interoperability testing

`lperf -c/-s`

Functions need to be verified
Handshaking (1st and 2nd sub-flow)
MP_CAPABLE
MP_KEY
MP_JOIN
MP_HMAC
MP_SEQ
MP_RTT
MP_CLOSE
MP_PRIO
MP_ADDADDR
MP_CONFIRM
MP_REMOVEADDR
Fallback mechanism

Interoperability test

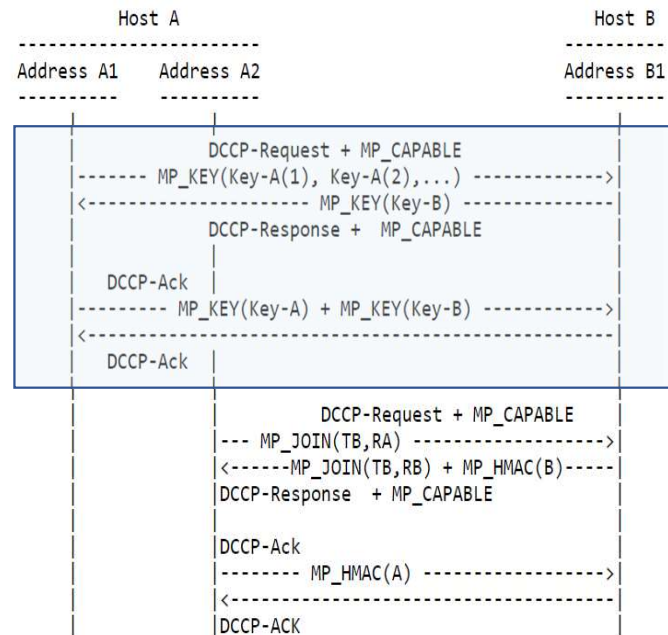
Function verified

Handshaking
1st subflow

MP_CAPABLE

MP_KEY

Draft specification



No.	Time	Source	Destination	Protocol	Length	Info
1	11:32...	1.20...	213...	DCCP	98	17665 → 5201 [Request] Seq=0 (serv
2	11:32...	213...	1.20...	DCCP	138	5201 → 17665 [Response] Seq=0 (Ack
3	11:32...	1.20...	213...	DCCP	110	17665 → 5201 [Ack] Seq=1 (Ack=0)
4	11:32...	213...	1.20...	DCCP	58	5201 → 17665 [Ack] Seq=1 (Ack=1)
5	11:32...	1.20...	213...	DCCP	98	62110 → 5201 [Request] Seq=0 (serv
6	11:32...	213...	1.20...	DCCP	162	5201 → 62110 [Response] Seq=0 (Ack
7	11:32...	1.20...	213...	DCCP	107	17665 → 5201 [DataAck] Seq=2 (Ack=
8	11:32...	1.20...	213...	DCCP	110	62110 → 5201 [Ack] Seq=1 (Ack=0)
9	11:32...	213...	1.20...	DCCP	58	52
10	11:32...	1.20...	213...	DCCP	1370	17665 → 5201 [Ack] Seq=3 (Ack=

4-Way handshake

Frame 1: 98 bytes on wire (784 bits), 98 bytes captured (784 bits)

Ethernet II, Src: JuniperN_f4:37:c2 (54:4b:8c:f4:37:c2), Dst: Dell_27:

Internet Protocol Version 4, Src: 1.202.162.49, Dst: 213.239.223.124

Datagram Congestion Control Protocol, Src Port: 17665, Dst Port: 5201

Source Port: 17665

Destination Port: 5201

[Stream index: 0]

Data Offset: 16

CCVal: 4

Checksum Coverage: 0

Checksum: 0x2ad1 [unverified]

[Checksum Status: Unverified]

Type: Request (0)

Extended Sequence Numbers: True

Sequence Number: 0 (relative sequence number)

Sequence Number (raw): 171616316566003

Service Code: not specified (0)

Options: (44 bytes)

Option Type: Multipath (46)

MP_KEY

Key Type: 0

Key Data: 0x3b7f9a7fbafde05b

Option Type: Timestamp (41)

Option Type: Change L (32)

Option Type: Change R (34)

Option Type: Mandatory (1)

Option Type: Change L (32)

Option Type: Mandatory (1)

Option Type: Change L (32)

Option Type: Change R (34)

MP_CAPABLE

version: 0

Interoperability test

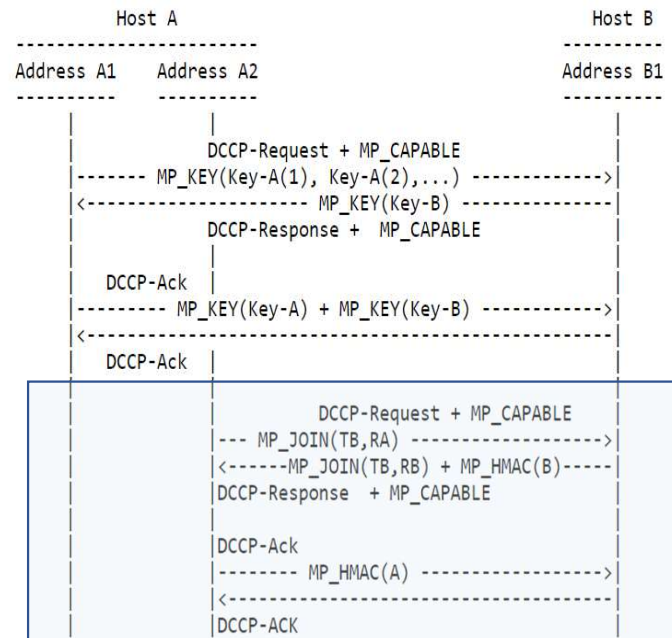
Function verified

Handshaking
2nd subflow

MP_JOIN

MP_HMAC

Draft specification



```
1 11:32:1... 1.202...213.239... DCCP 98 17665 → 5201 [Request] Seq=0 (service=
2 11:32:1... 213.23...1.202.16... DCCP 138 5201 → 17665 [Response] Seq=0 (Ack=0)
3 11:32:1... 1.202...213.239... DCCP 110 17665 → 5201 [Ack] Seq=1 (Ack=0)
4 11:32:1... 213.23...1.202.16... DCCP 58 5201 → 17665 [Ack] Seq=1 (Ack=1)
5 11:32:1... 1.202...213.239... DCCP 98 62110 → 5201 [Request] Seq=0 (service=
6 11:32:1... 213.23...1.202.16... DCCP 162 5201 → 62110 [Response] Seq=0 (Ack=0)
7 11:32:1... 1.202...213.239... DCCP 107 17665 → 5201 [DataAck] Seq=2 (Ack=1)
8 11:32:1... 1.202...213.239... DCCP 110 62110 → 5201 [Ack] Seq=1 (Ack=0)
9 11:32:1... 213.23...1.202.16... DCCP 58 5201 → 62110 [Ack] Seq=1 (Ack=1)
10 11:32:1... 1.202...213.239... DCCP 1370 17665 → 5201 [DataAck] Seq=3 (Ack=1)
11 11:32:1... 213.23...1.202.16... DCCP 62 5201 → 17665 [Ack] Seq=2 (Ack=3)
12 11:32:1... 1.202...213.239... DCCP 1370 17665 → 5201 [DataAck] Seq=4 (Ack=1)
```

Destination Port: 62110
[Stream index: 1]
Data Offset: 32
CCVal: 8
Checksum Coverage: 0
Checksum: 0x322f [unverified]
[Checksum Status: Unverified]
Type: Response (1)
Extended Sequence Numbers: True
Sequence Number: 0 (relative sequence number)
Sequence Number (raw): 20454788048994
Acknowledgement Number: 0
Acknowledgement Number (raw): 178992640069596
Service Code: not specified (0)

Options: (100 bytes)

Option Type: Padding (0)

Option Type: Multipath (46)

MP_JOIN: 1

Address ID: 0

Path Token: 0x85bae0ad

Nonce: 0x3f67591c

Option Type: Multipath (46)

MP_HMAC: 5

HMAC-SHA256: 87801cae9660519b70f6b96f3f06371fea7e5711

Option Type: Padding (0)

Option Type: Timestamp Echo (42)

4-Way handshake

MP_JOIN

MP_HMAC

Interoperability test

Function verified

MP_SEQ

MP_RTT

Draft specification

MP_SEQ

```

      1         2         3         4         5
01234567 89012345 67890123 45678901 23456789 01234567 89012345
+-----+-----+-----+-----+-----+
|00101110|00001001|00000100| Multipath Sequence Number
+-----+-----+-----+-----+
|
+-----+-----+
Type=46 Length=9 MP_OPT=4
```

MP_RTT

```

      1         2         3         4         5
01234567 89012345 67890123 45678901 23456789 01234567 89012345
+-----+-----+-----+-----+-----+
|00101110|00001100|00000110|RTT Type| RTT
+-----+-----+-----+-----+-----+
|           | Age                               |
+-----+-----+-----+-----+
Type=46 Length=12 MP_OPT=6
```

```

22 11:32:... 213.23...1.202.16... DCCP      70 5201 → 17665 [
23 11:32:... 1.202...213.239... DCCP    1394 17665 → 5201 [
24 11:32:... 213.23...1.202.16... DCCP      70 5201 → 17665 [
25 11:32:... 1.202...213.239... DCCP    1394 17665 → 5201 [
> Frame 23: 1394 bytes on wire (11152 bits), 1394 bytes captured
> Ethernet II, Src: JuniperN_f4:37:c2 (54:4b:8c:f4:37:c2),
> Internet Protocol Version 4, Src: 1.202.162.49, Dst: 213.239.1.202
> Datagram Congestion Control Protocol, Src Port: 17665, Dst Port: 5201
  Source Port: 17665
  Destination Port: 5201
  [Stream index: 0]
  Data Offset: 15
  CCVal: 5
  Checksum Coverage: 0
  Checksum: 0x3e35 [unverified]
  [Checksum Status: Unverified]
  Type: DataAck (4)
  Extended Sequence Numbers: True
  Sequence Number: 10 (relative sequence number)
  Sequence Number (raw): 171616316566013
  Acknowledgement Number: 3 (relative acknowledgement number)
  Acknowledgement Number (raw): 269606617297306
  Options: (36 bytes)
    ▾ Option Type: Padding (0)
      Padding: 00
    ▾ Option Type: Padding (0)
      Padding: 00
    ▾ Option Type: Padding (0)
      Padding: 00
    ▾ Option Type: Multipath (46)
      ▾ MP_RTT: 6
        RTT_Type: 0
        RTT: 604
        Age: 0
    ▾ Option Type: Multipath (46)
      Sequence Number: 9
    ▾ Option Type: Ack vector [Nonce 0] (38)
      Ack Vector [Nonce 0]: 00
```

⇒ MP_RTT

⇒ MP_SEQ

Interoperability test

Function verified

MP_CLOSE

Draft specification

Host A

DCCP-Close + MP_CLOSE
[B's key] [on all subflows]

Host B

<- Optional DCCP-CloseReq +
MP_CLOSE [A's key]
[on all subflows]

->

<- DCCP-Reset
[on all subflows]

```
467 11:32:... 1.202... 213.239... DCCP 1382 17665 → 5201 [DataAck] Seq=25
468 11:32:... 1.202... 213.239... DCCP 1382 17665 → 5201 [DataAck] Seq=25
469 11:32:... 1.202... 213.239... DCCP 74 17665 → 5201 [Close] Seq=25
470 11:32:... 213.23... 1.202.16... DCCP 66 5201 → 17665 [Reset] Seq=228
471 11:32:... 1.202... 213.239... DCCP 74 62110 → 5201 [Close] Seq=2 (A
472 11:32:... 213.23... 1.202.16... DCCP 62 5201 → 62110 [Reset] Seq=3 (A

Frame 469: 74 bytes on wire (592 bits), 74 bytes captured (592 bits)
Ethernet II, Src: JuniperN_f4:37:c2 (54:4b:8c:f4:37:c2), Dst: Dell_27:4
Internet Protocol Version 4, Src: 1.202.162.40, Dst: 213.239.223.124
Datagram Congestion Control
  Source Port: 17665
  Destination Port: 5201
  [Stream index: 0]
  Data Offset: 10
  CCVal: 2
  Checksum Coverage: 8
  Checksum: 0x2ce6 [unverified]
  [Checksum Status: Unverified]
  Type: Close (6)
  Extended Sequence Numbers: True
  Sequence Number: 257 (relative sequence number)
  Sequence Number (raw): 171616316566260
  Acknowledgement Number: 227 (relative acknowledgement number)
  Acknowledgement Number (raw): 269606617297530
  Options: (16 bytes)
    Option Type: Padding (0)
    Option Type: Padding (0)
    Option Type: Multipath (46)
      MP_CLOSE: 10
      Key Data: 0x1b9e36cff4447e05
    Option Type: Ack Vector [Nonce 0] (38)
```

Triggers the closure of all subflows

MP_CLOSE

Interoperability test

Function verified

MP_Prio

MP_ADDADDR

MP_CONFIRM

Draft specification

MP_Prio

										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
-----										-----										-----										-----									
0 0 1 0 1 1 1 0										0 0 0 0 0 1 0 0										0 0 0 0 1 0 0 1										Address ID									
-----										-----										-----										-----									
Type=46										Length=4										MP_OPT=9																			

MP_ADDADDR

1										2										3											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
0 0 1 0 1 1 1 0										var										0 0 0 0 0 1 1 1 Address ID											
Address (IPv4 - 4 bytes / IPv6 - 16 bytes)																															
Port (2 bytes, optional)															+ MP_HMAC option																
Type=46										Length										MP_OPT=7											

MP_CONFIRM

1	2	3	4	5
01234567 89012345	67890123 45678901	23456789 01234567	89012345	
00101110	var	00000000	List of confirmations ...	
Type=46	Length	MP_OPT=0		

1 17:07:...	80.187...	213.239...	DCCP	98 29987 → 5203	[Reque
2 17:07:...	213.23...	80.187.1...	DCCP	138 5203 → 29987	[Respc
3 17:07:...	80.187...	213.239...	DCCP	110 29987 → 5203	[Ack]
4 17:07:...	213.23...	80.187.1...	DCCP	58 5203 → 29987	[Ack]
5 17:07:...	80.187...	213.239...	DCCP	66 29987 → 5203	[Data]
6 17:07:...	213.23...	80.187.1...	DCCP	78 5203 → 29987	[Data]
7 17:07:...	80.187...	213.239...	DCCP	88 29987 → 5203	[Reque

```

  ▾ Options: (16 bytes)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Multipath (46)
  ▾ Option Type: Multipath (46)
    ▾ MP PRIO: 9
      MP_PRIO value: 5

```

⇒ MP_Prio

117:07:...	80.187...213.239... DCCP	98 29987 → 5203	[Request]
217:07:...	213.23...80.187.1... DCCP	138 5203 → 29987	[Response]
317:07:...	80.187...213.239... DCCP	110 29987 → 5203	[Ack] S
417:07:...	213.23...80.187.1... DCCP	58 5203 → 29987	[Ack] S
517:07:...	80.187...213.239... DCCP	66 29987 → 5203	[Data]
617:07:...	213.23...80.187.1... DCCP	78 5203 → 29987	[Data]
717:07:...	80.187...213.239... DCCP	88 29987 → 5203	[Request]

```

  ▾ Options: (28 bytes)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Padding (0)
    ▸ Option Type: Multipath (46)
    ▾ Option Type: Multipath (46)
      MP_CONFIRM: 0
    ▸ Option Type: Reserved (5)

```

⇒ MP CONFIRM

```

984 17:07:... 213.23... 80.187.1... DCCP      62 5203 → 29987 [Ack] Seq=4
985 17:07:... 213.23... 80.187.1... DCCP      90 5203 → 29987 [Data] Seq=
986 17:07:... 80.187.1... 213.23... DCCP      1282 29987 → 5203 [DataAck] S

```

```

  ▾ Options: (40 bytes)
    ▸ Option Type: Multipath (46)
    ▾ Option Type: Multipath (46)
      ▾ MP_ADDADDR: 7
        Address ID: 3
        Address: 10.110.10.2
      ▾ Option Type: Multipath (46)
        ▾ MP_HMAC: 5
          HMAC-SHA256: ad86d959c1

```

➡ MP_ADDADDR

MP_HMAC

HMAC-SHA256: ad86d959cbdbec540b8dd1b080ee8d81913959ef

Interoperability test

Function verified

MP_REMOVEADDR

Draft specification

MP_REMOVEADDR

```
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| 0 0 1 0 1 1 1 0 | 0 0 0 0 0 1 0 0 | 0 0 0 0 1 0 0 0 | Address ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
Type=46          Length=4          MP_OPT=8
```

Function verified

Fallback mechanism

Scenario: MP-DCCP is disabled immediately after 1st subflow establishment:

Attempt of establishing second subflow fails with DCCP-Reset

```
895 17:19:... 80.187... 213.239... DCCP 86 10023 → 5203 [Data] S
896 17:19:... 213.23... 80.187.1... DCCP 78 5203 → 10023 [Data] S
Options: (36 bytes)
  Option Type: Multipath (46)
  Option Type: Multipath (46)
    MP_REMOVEADDR: 8
    Address ID: 2
  Option Type: Multipath (46)
    MP_HMAC: 5
    HMAC-SHA256: 9fff6b3a15d2b1d969ee7526b3af48acd846c3d9
```

MP_HMAC

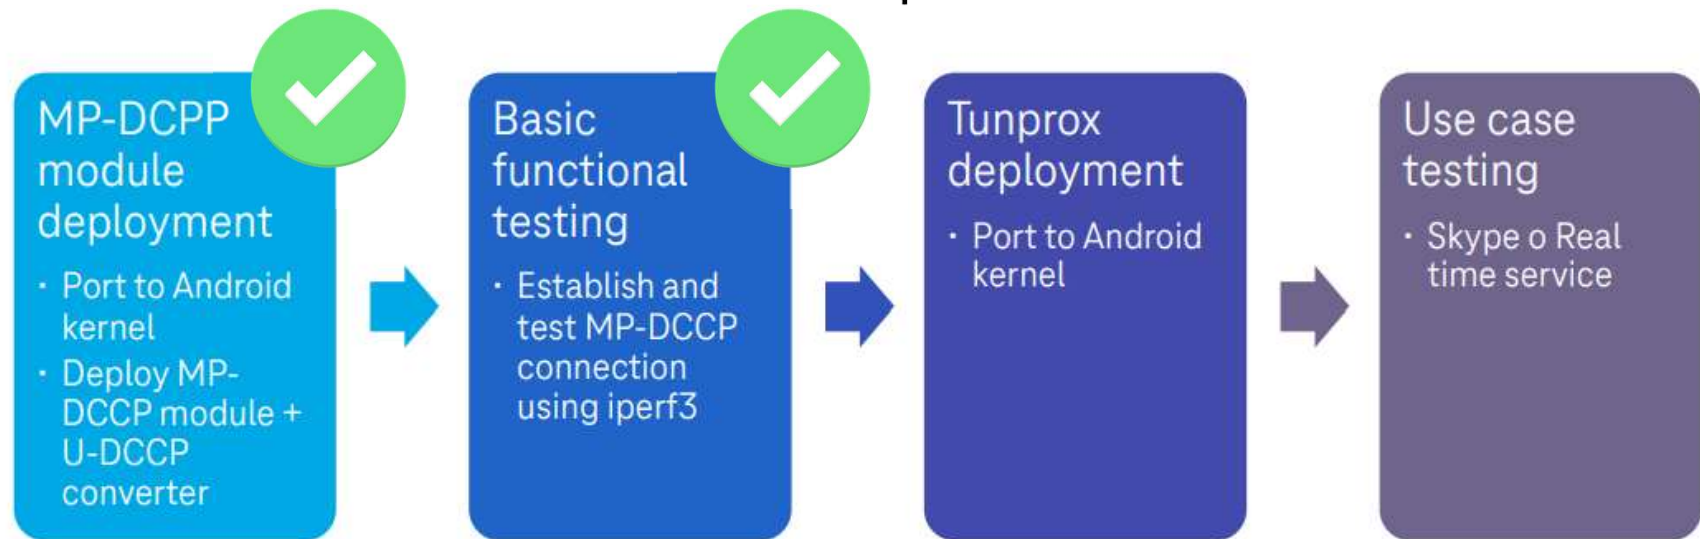
No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	80.18...	213.239...	DCCP	98	22397 → 5203 [Request] Seq=0 (servi
2	0.000001	213.2...	80.187.1...	DCCP	138	5203 → 22397 [Response] Seq=0 (Ack=
3	0.000002	80.18...	213.239...	DCCP	110	22397 → 5203 [Ack] Seq=1 (Ack=0)
4	0.000003	213.2...	80.187.1...	DCCP	58	5203 → 22397 [Ack] Seq=1 (Ack=1)
5	0.000004	80.18...	213.239...	DCCP	98	22399 → 5203 [Request] Seq=0 (servi
6	0.000005	213.2...	80.187.1...	DCCP	162	5203 → 22399 [Response] Seq=0 (Ack=0)
7	0.000006	80.18...	213.239...	DCCP	107	22397 → 5203 [DataAck] Seq=2 (Ack=1)
8	0.000007	80.18...	213.239...	DCCP	110	22399 → 5203 [Ack] Seq=1 (Ack=0)
9	0.000008	213.2...	80.187.1...	DCCP	62	5203 → 22399 [Reset] Seq=1 (Ack=1)
10	0.000009	80.18...	213.239...	DCCP	1370	22397 → 5203 [DataAck] Seq=3 (Ack=1)
11	0.000010	213.2...	80.187.1...	DCCP	62	5203 → 22397 [Ack] Seq=2 (Ack=3)
12	0.000011	80.18...	213.239...	DCCP	1370	22397 → 5203 [DataAck] Seq=3 (Ack=1)

> Frame 9: 62 bytes on wire (496 bits), 62 bytes captured (496 bits)
> Ethernet II, Src: Dell_27:49:ba (e4:43:4b:27:49:ba), Dst: JuniperN_f4:37:c2 (54
> Internet Protocol Version 4, Src: 213.239.223.124, Dst: 80.187.100.111
> Datagram Congestion Control Protocol, Src Port: 5203, Dst Port: 22399 [Reset] S

1st Subflow established

2nd Subflow reset

Next steps



- Deploy encapsulation framework to allow multipath transport of any IP traffic -> <https://datatracker.ietf.org/doc/html/draft-amend-tsvwg-multipath-framework-mpdccp-01> -> available at <https://github.com/telekom/tunprox>
- Realtime service tests, e.g. Skype

Conclusion

- The functions of MP-DCCP works well

Different combinations of MP-DCCP implementations are used for interoperability tests: 4.14 and 5.10 Linux reference implementation + Android 13 port (Xiaomi modified)

- Both LAN and WAN tests prove the completeness and maturity of MP-DCCP draft -06
- The Usage and operations are similar to MPTCP

Backup – Porting details

Port to MP-DCCP & U-DCCP converter

MP-DCCP and U-DCCP are both kernel mode programs. After adding their code to the kernel, we can compile the Android kernel with MP-DCCP function by opening the MP-DCCP related configuration options.

- MP-DCCP and U-DCCP repository

<https://github.com/telekom/mp-dccp>

<https://github.com/telekom/u-dccp>

- Add configuration items in gki_defconfig

Contains all functions of MP-DCCP prototype:

Scheduler selection

Path manager selection

Reordering engine selection

Path priority configuration

Congestion Control selection and queue setup

```
CONFIG_IP_DCCP=y
CONFIG_IP_DCCP_CCID2_DEBUG=y
CONFIG_IP_DCCP_CCID3_DEBUG=y
CONFIG_IP_DCCP_CCID5_DEBUG=y
CONFIG_IP_DCCP_DEBUG=y
CONFIG_IP_MPDCCP=y
CONFIG_IP_MPDCCP_DEBUG=y
CONFIG_MPDCCP_SCHED_SRTT=y
CONFIG_MPDCCP_SCHED_ROUNDROBIN=y
CONFIG_MPDCCP_SCHED_REDUNDANT=y
CONFIG_MPDCCP_SCHED_OTIAS=y
CONFIG_MPDCCP_SCHED_CPF=y
CONFIG_MPDCCP_SCHED_HANDOVER=y
CONFIG_MPDCCP_REORDER_FIXED=y
CONFIG_MPDCCP_STATS=y
CONFIG_U_DCCP=y
```

Compilation problems

The reasons for compilation errors are syntax errors and errors caused by the difference between the Android and Linux kernels

- Variables declared but not used

warning: unused variable 'peeked'

- Type error

const struct in_ifaddr *ifa -> struct in_ifaddr *ifa

- Common data structure changes require both common and msm repositories to be changed

```
struct sk_buff {  
    ...  
    char cb[128] __aligned(8); //48->128  
};
```

- Different function names in Linux and Android kernels

flowi4_to_flowi_common->flowi4_to_flowi