An SDN and CPS Based Opportunistic Upload Splitting for Mobile Users

Maurizio Casoni, Carlo Augusto Grazia, and Martin Klapez(✉)

Department of Engineering "Enzo Ferrari", University of Modena and Reggio Emilia, via P. Vivarelli 10, 41125 Modena, Italy
{maurizio.casoni,carloaugusto.grazia,martin.klapez}@unimore.it

Abstract. This paper proposes an hybrid approach composed by Software Defined Networking (SDN) and Cyber-Physical Systems (CPS) to boost the upload speed of mobile users in low-bandwidth environments through a next generation Mobile Collaborative Community (MCC). The core idea is to use a high-bandwidth local communication system, like IEEE 802.11 (WiFi), in order to distribute data efficiently through mobile hosts; then, the distributed data may be sent from each mobile node to the original destination through their low-bandwidth mobile interface for wide area network communication. With our solution some drawbacks of MCC are faced. With the use of SDN we defined a flexible and easy-to-configure MCC system which operates in a transparent way for the end hosts. At the same time, the use of CPS creates a feedback for the system regarding the hosts channel status; this way the system is able to fully exploit the MCC potential by increasing the upload speed for both congested and non-congested scenarios. We demonstrate the efficiency of our solution through experimental results obtained using the Mininet network emulator where POX and a Pyretic controller serve as a dynamic data repartition engine.

Keywords: Collaborative networks · CPS · MCC · SDN

1 Introduction

In the last decades, the massive introduction of mobile nodes as well as novel mobile technologies started to steer research in the networking areas. Mobile users are gaining Internet access through a widespread number of wireless technologies, and wireless networks are becoming ubiquitous. This trend has led the industry to introduce technologies for mobile wireless data communication, such as GPRS, EDGE, UMTS, HSPA and LTE. Several research activities attempted to enhance network performance, focusing on the concurrent use of multiple wireless technologies available on a host [1–4].

One possible optimization comes from users collaboration; in fact, almost all the available smartphones can benefit from a high-speed Wireless LAN (WLAN) interface and a Wireless WAN (WWAN) link which may suffer more in terms of performance and stability than the "local" links. The idea of collaboration

B. Mandler et al. (Eds.) IoT 360° 2015, Part I, LNICST 169, pp. 67–76, 2016.
DOI: 10.1007/978-3-319-47063-4_6

is based on sharing the WWAN link bandwidths with other members of the collaborative network through the high-speed WLAN interface. This approach takes the name of Mobile Collaborative Community (MCC).

Different paradigms are also evolving so as to integrate novel technology and, following an opportunistic and collaborative approach, to provide benefits for end users. An example of this novel solutions are Cyber-Physical Systems (CPS) [5], sometimes also presented as cross-layering solutions for IoT [6]. In such systems the network nodes are exploited both as a computing terminal, as always, but also as sensors. Each node could provide information to a core system in order to maximize some functions like throughput, delay, reliability, security and so on. One of the main challenges of these systems is the complexity introduced in the network sustainability due to the presence of these sensors [7]

A different paradigm investigated in the last years is Software Define Networking (SDN), in which the network data plane is decoupled from its control plane. Following this approach the setup power of a network explodes, giving to specific nodes the possibility to filter and modify IP packet fields, to change the traffic path and to forward packets following a specific optimization function. Again, the optimization could be provided for different figures of merit like throughput, delay, load balancing [8] and so on; even business aspects have been managed through SDN solutions [9].

This two main research areas are starting to be connected to each other [10,11], and that is essentially the purpose of this paper. What we did in this work has been to merge together solutions coming from the SDN world with the CPS one, applied to the MCC problem. Throughout this solution we aim to sensibly improve the upload transfer time of a mobile user by exploiting the resources shared by the LAN neighbors.

The discussion is organized as follows: Sect. 2 summarizes related work. Section 3 describes our proposal. Section 4 introduces the test environment, while Sect. 5 shows the emulation results. In Sect. 6 the conclusions of our work are drawn.

2 Related Work

In this section we revise the literature about collaborative algorithms aimed at boosting network throughput.

Work based on MCCs have been already proposed [12–14] by several authors. In this papers, Ad-hoc MCCs have been proposed as a solution to address communication hurdles. MCC enables two or more persons to aggregate their low-bandwidth mobile network channels to achieve a virtual high-bandwidth channel for collaborative data transfer. At the same time nodes communicate with each other using their high-bandwidth local area network, as instance through IEEE 802.11 (WiFi) links. The main issue of the already proposed approaches is that they involve ISP networks and make use of algorithms that run among terrestrial infrastructures; moreover, they are not transparent for end users, needing host modification and particular scheduling algorithms as well. Last but not least, these works focus on improving downlink speeds, while the upload problem is never considered.

A different yet correlated work proposes some improvements to the MCC topic with PRISM [15], a proxy-based inverse multiplexer that enables TCP to efficiently utilize the community members. This work tries to solve the degradation performance of TCP communications that occur over MCC systems due to frequent out-of-order packet deliveries. Unfortunately, all the issues discussed before for general MCC solutions remain unsolved.

Another interesting work is a packet scheduler called DAPS [16], provided to mitigate TCP issues that incur when different interfaces are used to perform an end-to-end transfer, for example through Multipath TCP (MPTCP). DAPS introduces a smart dispatcher in charge of deciding the balance of traffic among the different interfaces, in order to maintain a linear order of arrival of packets and therefore avoiding the constriction of TCP sender window size. This idea has been revisited in our work, that instead implements an SDN controller to maintain a good balance between different interfaces, in our case represented by different nodes.

3 Solution Description

To introduce the significance of cooperative solutions we explore the context of ad-hoc networks deployed for emergency purposes [17]. We present a scenario that has been extracted from the EU FP7 project Public Protection and Disaster Relief - Transformation Center (PPDR-TC), where there are a certain number of First Responders who connect to a MEOC (Mobile Emergency Operations-Control Centre) that brings to them IEEE 802.11 coverage, thus representing a deployable (and mobile) WiFi common network for operators to exchange data. First Responders may belong to different entities, such as medical personnel, security services or firemen, and it is assumed that each one of them has at disposal a common smartphone/device connected to the respective mobile service. In unplanned disasters, however, terrestrial infrastructures are often damaged or congested, and therefore the capabilities offered by mobile networks are more often than not restricted at least. This scenario is depicted in Fig. 1. It should however be noticed that our proposal is applicable in all environments where the upload speed is constrained or should be maximized, as for example may be a domestic or office WLAN where connectivity suffers significant bandwidth drops or where it is hindered by legacy telephone lines (e.g. ADSL).

Returning to our reference scenario, suppose that a generic field operator wants to upload a file to a remote destination, an operation often needed or desired by PPDR personnel [18]; instead of using its low-bandwidth (or congested) link, this host sends first the file(s) to the MEOC using the high-quality LAN channel. The MEOC then acts as an OpenFlow switch, with an SDN Controller in charge of dynamically dispatching the incoming packets to the other field operators, which in turn *concurrently* forward the packet to the remote destination. Here is where CPS capabilities come into play; assumed that the mobile

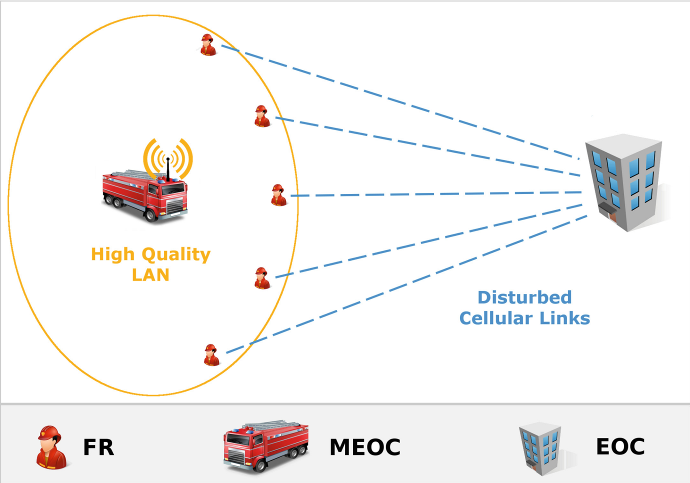

Fig. 1. Reference scenario

link type of a generic field operator is not known in advance to the MEOC, and that operators may belong to different services with different mobile contracts, field hosts may tell the MEOC the bandwidth they have available from the cellular channels with a single packet. Common smartphones have already embedded the capability of automatically detecting the signal quality of a channel; therefore, if the link degrades or, on the contrary, if the link improves (e.g. passing from 2G to 3G), field hosts may communicate the alteration to the MEOC with a simple application that runs in the smartphone user-space. CPS capabilities, however, prove even more useful in trickier cases, as when the signal coverage remains good but performance are degraded by congestion. The application can easily perform a throughput test at pre-determined intervals, and then send a packet to inform the MEOC SDN Controller of the current link capabilities.

4 Emulation Environment

To test the Cyber-Physical Mobile Collaborative Community detailed in the previous Section, we used the Mininet network emulator version 2.2.0 running on a VirtualBox Ubuntu 14.04 Virtual Machine. To manage SDN operations, we used OpenFlow 1.0 as Southbound API, while Pyretic based on POX 0.2.0 has been used to write the controller application that constitutes the Northbound API.

The architecture of the test system is depicted in Fig. 2. The MEOC is composed by a host (inserted only for consistency with the previous discussion, as the MEOC host is ignored in our tests) and an OpenFlow switch. There are five client hosts (representing field operators), each with its own LAN connection to the MEOC and with a specific mobile connection to a remote node. The latters are modeled as specified in Table 1.

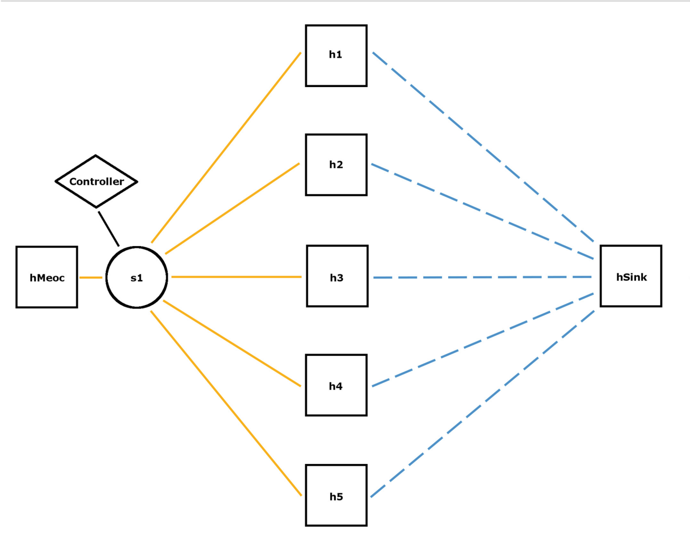

Fig. 2. Test system architecture

Table 1. Links configuration

Client host	Radio technology simulated	Max. upload speed	Latency
h1	EDGE	118 kbit/s	100 ms
h2	EDGE	118 kbit/s	100 ms
h3	UMTS	128 kbit/s	70 ms
h4	HSPA Rel.5 (HSDPA)	384 kbit/s	45 ms
h5	EDGE	118 kbit/s	100 ms

The values in the table have been extracted from [19] and from direct measures conducted by the authors. The LAN channels between client hosts and the MEOC are assumed to be high-quality, and therefore no change to the default Mininet link setup has been made. Because the clients mobile radio channels are operating at low-bandwidth, no change to the default Mininet link setup regarding queue sizes has been made either. In each client host, a Linux Bridge has been configured to forward traffic properly between its two interfaces, and the same holds for the remote node, where all its five interfaces have been bridged; both the "brctl" and "bridge" tools have been used to configure the Linux Bridges.

To measure the upload throughput we used Iperf 2.0.5, running a default Iperf server on the remote node and a default Iperf client on a specific client host (*h1* of Table 1), therefore generating regular TCP traffic. We have chosen to let a host having an EDGE link to be the one that needs uploading data because 2.5G is usually the fallback technology in rural areas or when faster solutions are not available anymore. As Table 1 attests, we chose not to artificially "pump" our results by introducing significantly faster channels in great number. The data rate is calculated on the remote host (*hSink* in Fig. 2).

All tests belong to one of the following cases:

1. *h1* sends data directly to the remote host using its EDGE channel.
2. *h1* sends data to MEOC first, then the MEOC forwards the packets coming from *h1* to hosts *h2*, *h3*, *h4*, *h5* using a Round-Robin (RR) discipline.
3. *h1* sends data to MEOC first, then the MEOC forwards the packets coming from *h1* to hosts *h2*, *h3*, *h4*, *h5* using a Weighted Round-Robin (WRR) discipline, calculated on the basis of CPS information.

Specifically, in the latter case for each packet assigned to a node with an EDGE link the Controller assigns two packets to UMTS links and three to HSDPA links, thus serving the nodes in a Weighted Round-Robin fashion thanks to the knowledge provided by the CPS engine of our proposal. In the latter two cases, the channel between *h1* and the remote host is kept free for TCP Acknowledgement packets.

Last but not least, for each of the aforementioned cases, four different tests have been carried out in order to assess the system performance under different network conditions, i.e. with no or low losses (0 % and 2 % of packet losses, respectively) and with high or very high losses (5 % and 10 % of packet losses, respectively). This brings the total tests configurations to twelve.

The complete testbed with instructions about how to perform these tests is available at [20]. The Section that follows presents tests results.

5 Performance Analysis

First of all, let us discuss about the rationale behind our results selection. Out of a total of 12 test configurations, we performed 10 runs for each of them. As representative of the results, we picked up the average value. The variance itself has not been plotted because the tests were run on a Virtual Machine (VM) (see Sect. 4); unfortunately, due to the nature of VMs, interferences from processes running on the VM host operating system cause the CPU allocation to the VM itself to drop and oscillate. Therefore, we cannot assess at this point if the registered variance is due to this reason or some other, and is thus pointless to plot the relative figures.

Figure 3 reports the results of the first test suite, when no losses occur. The upload data rate of the default configuration (i.e. without any Controller) is equal to 113 kbit/s, a little less than the maximum EDGE throughput which is of 118 kbit/s. With MCC RR, the data rate attests itself to 396 kbit/s, for a 3.5X Speedup, while with MCC WRR the same figure is equal to 424 kbit/s, for a 3.7X Speedup. Please note that a linear growth of network throughput is hindered by the out-of-order delivery of packets.

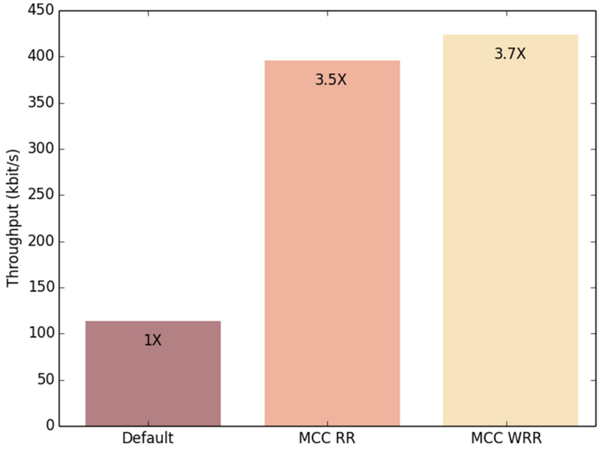

Fig. 3. Upload throughput with no losses

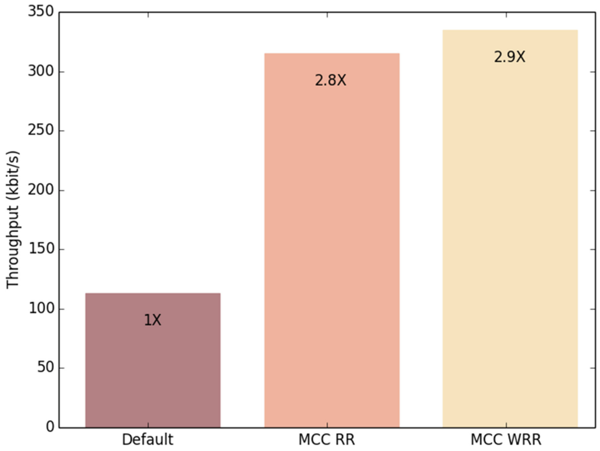

Fig. 4. Upload throughput with 2 % of packet loss

Figure 4 shows the test results with 2 % of packet losses. In this case, with MCC RR the data rate is equal to 315 kbit/s for a 2.8X Speedup, while with MCC WRR the throughput attests to 335 kbit/s, for a 2.9X Speedup.

The improvement given by our proposal starts to decrease significantly when losses are high (i.e. 5 % of packet losses), as Fig. 5 testifies. Here, with MCC RR the throughput is equal to 192 kbit/s, for a 1.7X Speedup, while with MCC WRR the data rate attests to 225 kbit/s for a 2X Speedup.

This behavior is more prominent in case of very high losses (i.e. 10 % of packet losses). Figure 6 shows how with MCC RR the upload speed reaches 140 kbit/s for a 1.2X Speedup, while MCC WRR allows 147 kbit/s for a 1.3X Speedup.

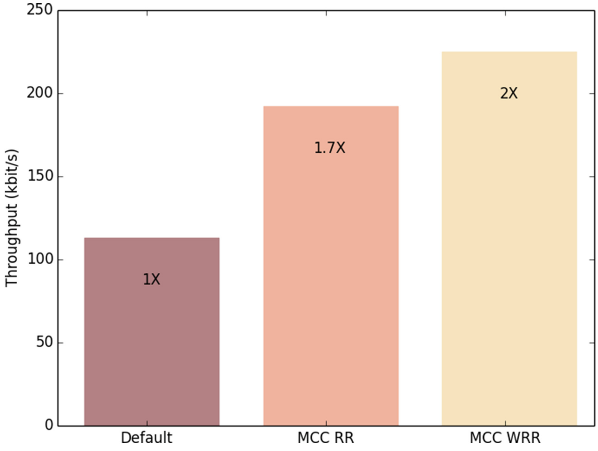

Fig. 5. Upload throughput with 5 % of packet loss

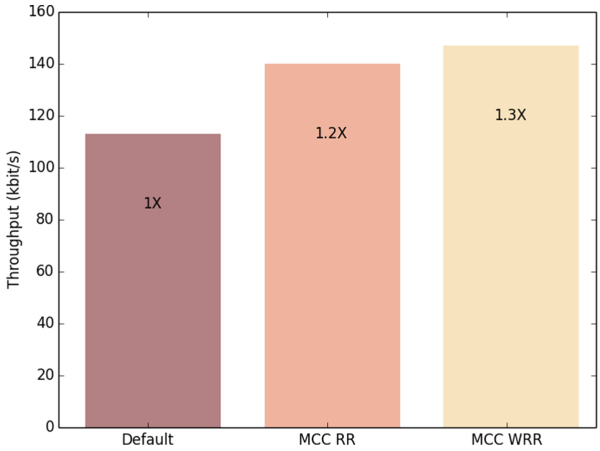

Fig. 6. Upload throughput with 10 % of packet loss

It can be therefore concluded that, even in cases of very high packet losses, our proposal can benefit the end users, although not with the same performance variance seen in more common cases.

In the case of frequent losses, it is not only the out-of-order delivery that hinders the performance growth, but also the necessary TCP retransmissions.

6 Conclusions

In this work, we proposed a Cyber-Physical Mobile Collaborative Community system to enhance upload speeds for client hosts in constrained scenarios.

The proposal is based on an SDN Controller that serves as a dynamic packet dispatcher to neighbor client nodes, that in turn are able to concurrently forward the original packets to their proper destination. This way, it is possible to realize a virtual multipath connection with a solution that is completely transparent for client hosts. Emulation results performed with Mininet confirm that the performance improvement is significant. Furthermore, they point out how the exploitation of CPS capabilities is able to further improve throughput, in a measure that varies from +5 % to +15 % with respect to the MCC system alone.

Acknowledgments. This work was also supported by the European Commission under PPDR-TC, a collaborative project part of the Seventh Framework Programme for research, technological development and demonstration. The authors would like to thank all partners within PPDR-TC for their cooperation and valuable contribution.

References

1. Han, H., Shakkottai, S., Hollot, C.V., Srikant, R., Towsley, D.: Multi-path tcp: a joint congestion control and routing scheme to exploit path diversity in the internet. IEEE/ACM Trans. Netw. **14**(6), 1260–1271 (2006). doi:10.1109/TNET.2006.886738
2. Carter, C., Kravets, R.: User devices cooperating to support resource aggregation. In: Proceedings Fourth IEEE Workshop on Mobile Computing Systems and Applications 2002, pp. 59–69 (2002)
3. Hsieh, H.-Y., Sivakumar, R.: A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In: Proceedings of the 8th Annual International Conference on Mobile Computing and Networking, ser., MobiCom 2002, pp. 83–94. ACM, New York (2002). http://doi.acm.org/10.1145/570645.570656
4. Rodriguez, P., Chakravorty, R., Chesterfield, J., Pratt, I., Banerjee, S.: Mar: a commuter router infrastructure for the mobile internet. In: Proceedings of the 2nd International Conference on Mobile Systems, Applications, and Services, ser. MobiSys 2004, pp. 217–230. ACM, New York (2004). http://doi.acm.org/10.1145/990064.990091
5. Park, K.-J., Zheng, R., Liu, X.: Cyber-physical systems: milestones and research challenges. Comput. Commun. **36**(1), 1–7 (2012). http://www.sciencedirect.com/science/article/pii/S0140366412003180
6. Han, C., Jornet, J.M., Fadel, E., Akyildiz, I.F.: A cross-layer communication module for the internet of things. Comput. Netw. **57**(3), 622–633 (2013). http://www.sciencedirect.com/science/article/pii/S138912861200357X
7. Shin, D.-H., He, S., Zhang, J.: Robust and cost-effective design of cyber-physical systems: an optimal middleware deployment approach. IEEE/ACM Trans. Netw. **99**, 1–14 (2015)
8. Adami, D., Giordano, S., Pagano, M., Santinelli, N.: Class-based traffic recovery with load balancing in software-defined networks. In: Globecom Workshops (GC Wkshps), 2014, pp. 161–165 (2014)
9. Adami, D., Donatini, L., Foddis, G., Giordano, S., Roma, S., Topazzi, S.: Design and development of management functions for distributed monitoring based on sdn-based network. In: Euro Med Telco Conference (EMTC), 2014, pp. 1–5 (2014)

10. Skowyra, R.W., Lapets, A., Bestavros, A., Kfoury, A.: Verifiably-safe software-defined networks for cps. In: Proceedings of the 2nd ACM International Conference on High Confidence Networked Systems, ser. HiCoNS 2013, pp. 101–110. ACM, New York (2013). http://doi.acm.org/10.1145/2461446.2461461
11. Qin, Z., Do, N., Denker, G., Venkatasubramanian, N.: Software-defined cyber-physical multinetworks. In: 2014 International Conference on Computing, Networking and Communications (ICNC), pp. 322–326, February 2014
12. Lee, J., Choi, J.-G., Bahk, S.: Opportunistic downlink data delivery for mobile collaborative communities. Comput. Netw. **57**(7), 1644–1655 (2013)
13. Jang, I., Suh, D., Pack, S.: Minimizing content download time in mobile collaborative community. In: 2014 IEEE International Conference on Communications (ICC), pp. 2490–2495, June 2014
14. Kim, K.-H., Shin, K.G.: Improving tcp performance over wireless networks with collaborative multi-homed mobile hosts. In: Proceedings of the 3rd International Conference on Mobile Systems, Applications, and Services, ser. MobiSys 2005, pp. 107–120. ACM, New York (2005). http://doi.acm.org/10.1145/1067170.1067183
15. Kim, K.-H., Shin, K.: Prism: improving the performance of inverse-multiplexed tcp in wireless networks. IEEE Trans. Mob. Comput. **6**(12), 1297–1312 (2007)
16. Kuhn, N., Lochin, E., Mifdaoui, A., Sarwar, G., Mehani, O., Boreli, R.: Daps: intelligent delay-aware packet scheduling for multipath transport. In: 2014 IEEE International Conference on Communications (ICC), pp. 1222–1227, June 2014
17. Casoni, M., Grazia, C.A., Klapez, M., Patriciello, N.: Integration of satellite and LTE for disaster recovery. IEEE Commun. Mag. **53**(3), 47–53 (2015). doi:10.1109/MCOM.2015.7060481
18. PPDR-TC Consortium: PPDR's Needs and Requirements. Project Deliverable D2.2, January 2014
19. Svoboda, P., Ricciato, F., Keim, W., Rupp, M.: Measured WEB performance in GPRS, EDGE, UMTS and HSDPA with and without caching. In: IEEE World of Wireless, Mobile and Multimedia Networks, 2007 (WoWMoM 2007), June 2007
20. Testbed with instructions: netlab.ing.unimo.it/sw/cpmcc.zip