SCIENCE AND ENGINEERING RESEARCH COUNCIL
RUTHERFORD APPLETON LABORATORY

COMPUTING DIVISION

DISTRIBUTED INTERACTIVE COMPUTING NOTE 690

COMMON BASE COMMUNICATIONS
Report Sept 82

Issued by
W P Sharpe

28 September 1982

DISTRIBUTION:
R W Witty
D A Duce
K Robinson
A S Williams
C Prosser
L O Ford
J M Loveluck
K J Fermor
W P Sharpe
CBP/Networks/General File

1. INTRODUCTION

The purpose of this paper is to summarise progress to date in provision of communication facilities for the common base program and to define the work plan for the next quarter.

2. CURRENT POSITION

2.1 PDP11

2.1.1 PDP11 BBP

CPW has completed and distributed the implementation of BBP in the Logica UMCZ80 front end together with a host driver for Unix V7. The driver is completely configurable for byte ordering and basic block type options to meet the old and new standards. The distribution includes Z80 programming tools including the Kent BCPL cross compiler and cross assembler; the latter have been tidied up for easy installation.

Logica VTS now have firmware for the UMCZ80 implementing BBP. The PROMs containing the firmware can be fitted to existing boards and will be supplied free for all the UMCZ80 boards we have, if and when we want them. THe Logica firmware is not configurable for byte order and block type and the host interface is of course different to the RAL driver so we cannot make immediate use of it.

2.1.2 PDP11 BSP

A version of BSP for file transfer is available. An FTP daemon process has been written by WPS that sets up transfers using SSP (Single Shot Protocol). Transfers are initiated by a user program with the same interface as PUFTP. This software is definitely `interim' not `production' code.

2.1.3 PDP11 TSBSP

KJF has substantially written the TSBSP module according to the design of DCS 585. Limited testing has been performed by simulation of external events. The definition of TSBSP is that of the Kent August 82 document. Work has not started on the other modules of the implementation (user interface routines, and kernel link module).

2.1.4 PDP11 X25, XX, TS29, FTP80, JNT Mail

Good progress has been made by KJF in testing the York X25 front end and the higher level protocols. At the last Unix networking meeting Keith Ruttle (York) was confident that a stable release (Version 1.1) could be reached within a few weeks. Quality assurance has proved to be a problem because of the lack of written definitions of legal network behaviour on SERCNET.

2.1.5 PDP11 MACE

The MACE is a 6809 based front end processor for implementing the the protocols up to transport service over the ring (as the York LSI11 does for WANs). The MACE also allows direct access to the lower levels of protocol for user processes to implement their own transport protocols (RPC etc). The MACE is being developed by Orbis and ICL (the old Dataskil part) under a JNT contract, together with host attachment hardware and software for PDP11 Unix, VAX Unix, VAX VMS, Prime, GEC, 68000. For Unix the user program interface to the transport level will be identical to that of the York WAN package, allowing porting without tears of the high level protocols (TS29, FTP80, Mail). The MACE will be available 3Q83 and probably be the first extant implementation of the standard ring protocols (see 2.3 below). Unit price for a MACE, complete with ring node is currently set at $2500.

2.2 Perq

2.2.1 Perq BBP & BSP

DIC Note 660 describes the implementation by JML of the Kent BBP and BSP package on the Perq. This package is used for file transfer and has been given the same user interface as PUFTP. The GPIB ring interface is used subject to restrictions imposed by the code in the Z80 board. The current working implementation is under POS; BBP has run under Unix and no problems are expected in mounting the rest. Note that new standard P-Service has not been implemented (ie only type 0 Basic Blocks have been used). The same software has also been adapted to run on serial lines between Perqs and between a Perq and the PDP11/70.

2.2.2 Perq X25, XXX, FTP80 etc

A contract has been placed with York to port the LSI11 front end to the Perq with all the high level protocols. The contract also includes a study for a long term solution for Perq X25. York will explore the possibility of using the new Falcon board from DEC to engineer a single LINK + LSI11 board that could be mounted in the Perq.

ERCC also have a contract to produce a proposal for Perq X25. They are exploring use of the SIO port on the Z80 board with complete inboard implementation of X25.

2.2.3 Perq MACE

ICL will shortly be providing a proposal for putting the MACE on the Perq. CAMTEC, who are making the JNT ring terminal concentrator, are producing a modular architecture which could also be the basis of a MACE type front end. They have indicated that they too would like to put in a proposal and have been informally invited to do so.

2.2.4 Perq GPIB ring interfaces

CAMTEC are to be given a contract to build 50 interfaces to the RAL design. Delivery is quoted as 12 weeks. DIC Note 660 gives details of current performance of this interface.

2.2.5 Perq FTP over Asynchronous Lines

This refers to the implementation of FTP77 over an asynchronous protocol for serial lines, invented at Atlas and implemented on Primes and GECs - not PUFTP. It has been decided as an interim measure to use this protocol as a means of distributing software for Perqs in the field i.e. software will be distributed over SERCNET and then down serial lines from ICF hosts. Sussex already have an implementation written in C for Unix on PDP11 and have agreed to port this to the Perq. RAL had a version of the PDP11 software that was not very reliable. The implementation on the Perq will require an RS232 driver to be written at RAL.

2.3 Ring Standards

2.3.1 Current Generation

The hardware specification is in the final stages of preparation for printing and will be published as "Cambridge Ring 82, Interface Specifications". The existing RAL and Logica equipment does not meet the standard in all respects, but only the mechanical aspects (connectors) are likely to prove any real inconvenience. It has been necessary to reverse the Polynet byte ordering conventions and modified ROMs to achieve this have been distributed to all DCS sites. The standard requires 40 bit minipaackets; all equipment installed in RAL is being converted, but there are no immediate plans to convert DCS equipment in the universities.

The protocol specifications are complete and should be available within a few weeks. They will be published as "Cambridge Ring 82, Protocol Specifications." The protocols are now named as follows:

BBP	P-Service	(Packet service)
BSP	V-Service	(Virtual Circuit Service)
TSBSP	N-Service	(Network service).

The V and N service differ from BSP and TSBSP in numerous small but significant ways. To avoid confusion therefore it is advisable to refer to the new protocols as CR82 V-Service etc.

These two specifications will join the JNT "rainbow series" as the"orange books".

2.3.2 Next Generation

Some effort was expended by WPS and P E Bryant in bringing together manufacturers interested in developing a standard for the next generation of slotted rings. Slotted ring technology scales up to higher speeds very readily and Cambridge University already have a prototype network running at 80 Mbits/sec. Such an initiative is supported by the FOCUS LAN report, but the activity has languished through lack of manufacturer commitment.

2.4 Distributed System Protocols

In parallel with the use of the virtual circuit based protocols (BSP, TSBSP) with FTP etc for the connection of heterogeneous systems there is widespread experimentation, and some experience, with the use of simpler protocols for many of the inter-process communication (IPC) and remote procedure call (RPC) requirements of more homogeneous distributed systems. Currently the situation can be regarded as chaotic with a number of examples showing that things can be done with `lightweight' protocols but very few recommendations on how such protocols should be structured.

The Newcastle Connection is based on a `lightweight RPC' and Newcastle were very critical of the orange book exercise since no effort was made to consider these alternatives. A working group has now been set up by WPS and PE Bryant with a view to understanding the principles of these protocols and producing some recommendations. The working group has held its first meeting. This activity is also driven by the need for some protocols to run on some of the RAL local networks being installed in other divisions.

2.5 Ring-X25 Gateway

There is still no sign of a satisfactory product appearing without some initiative and more money from JNT/SERC. Solutions based on the MACE or the CAMTEC product are being actively investigated; a functional requirements document has been circulated to the relevant manufacturers. There is reasonable hope that a gateway can be produced at about the same time as the MACE i.e. 3Q83.

3. DEVELOPMENT PLAN

3.1 General

This section reflects decisions taken at a meeting between RWW, KR, LOF, CP, DAD & WPS.

It has been decided to adopt a distributed architecture for CBP systems based upon the Newcastle Connection. This allows a number of Unix systems that are linked by a Ring to operate from a user point of view as a single system. The systems communicate one with another using the Newcastle Remote Procedure Call, which is implemented directly over the P-Service. The RPC implements Unix systems calls as remote procedures and therefore links all machines that run Unix i.e. it can link up Perqs and PDP11s etc. To satisfy the communications requirement it is sufficient for each Perq to be equipped with a cheap GPIB ring interface and to have the P-Service implemented in-board. It is worth stressing that the traditional network protocols are not relevant between systems that are linked in this way any more than they are relevant between two users on the `same' Unix machine.

When it is necessary to communicate with another network, or with heterogeneous systems on the same local network then a user program must have access to an implementation of the transport service over the WAN or LAN. The implementation need not reside on the same machine. In OSI parlance, the distributed Unix system is a single `end-system', the components of which communicate with one another by some private parochial means. It is therefore possible to provide a transport connection over the Ring and X25 in only one of the linked systems as a shared resource for the entire distributed system. It is then economically feasible to use the MACE to provide the N-Service since only one system need be equipped. Likewise, only one system will need to have a LSI11 X25 front end. Further, since the distributed Unix acts as a single end system the need for a specific gateway disappears - a gateway is only necessary where transport connections from one network must be relayed on to multiple end systems on another network.

3.2 Actions

(1) KJF will continue to work on the testing and debugging of the York package and will be responsible for configuring all the LSI11 systems before they are sent out. KJF has been invited to attend SIDM meetings, which will enable him to ensure that the York Unix package stays in line with SERC requirements. KJF and WPS will prepare a presentation of the facilities offered by the package to be given at a section meeting when release 1.1 has been mounted on the PDP11/70.

(2) 50 GPIB Ring interfaces will be ordered (WPS):- 20 for STI, 20 for CBP and 10 for DCS. JML will develop a suite of test programs for acceptance testing and fault diagnosis.

(3) Newcastle will be provided with Perqs for them to integrate the Newcastle RPC into Accent (KR).

(4) JML will supply a P-Service driver for Perq Unix that can be used by Newcastle.

(5) CPW will mount the Newcastle RPC onto the PDP11/34 and the PDP11/70 as soon as it is available.

(6) There are various aspects of the current implementation of the Newcastle RPC that are, or are likely to be unsatisfactory from the point of view of the CBP and general applicability. In particular:

(i) The ACCENT definitions for typed messages should be used in place of ad hoc conventions.

(ii) A name server should be used for locating machines, for ease of system configuration.

(iii) The addressing mechanism should integrate properly with the ACCENT IPC usage.

These and other issues must be considered by CPW and WPS in conjunction with Newcastle. CPW and WPS will continue with the lightweight protocols group with the aim of producing coherent definitions for the Newcastle RPC and other LAN services.

(7) Procurement of a MACE or other equivalent interface for the Perq must be set in hand ASAP (WPS).

(8) Current work on TSBSP is suspended but may be resumed on the Perq when resources are available. RWW and DAD will also discover whether Kent or elsewhere wish to produce a Perq implementation. This is seen only as a `fall back'.

(9) As a low priority the Logica BBP firmware will be installed a modified host driver produced to drive it (CPW).

(10) In order to allow Sussex to implement the FTP over asynchronous lines an RS232 driver for the Perq will be written as soon as possible (LOF). The Sussex software will be subjected to thorough testing at RAL before release.

(11) The Perq and PDP11 interim FTP facilities (2.1.1 & 2.2.1) will be adjusted to use a common start up procedure so that they can be used between Perqs and PDP11s. No effort will be expended to converting the Perq version to P-Service since the UMCZ80 can run BBP. This should ensure that it eventually goes out of service. (WPS, JML).

3.3 Other Notes

(1) No effort is allocated to work on Ethernet. However, the Newcastle RPC should be structured in such a way as to promote portability.

(2) CPW will be allocating some time to proposals for mounting LCF on Perq.

(3) There is no immediate requirement for a network software to implement the functional distribution of ACCENT over the ring.