CPA - Communicating Process Architectures
Observations / thoughts
A personal view

Barry M Cook

August 1, 2014

1 Introduction

CPA is the continuation of a long series of conferences related to a computer development in the 1980's.

It appears that this conference now has some difficulty attracting papers and, perhaps to a lesser extent, attendees.

I am wondering why this is. Could it be:

- The subject matter is dead?
- It is unclear what "the subject matter" is?
- There is no place for a CPA-style conference?
- Something else?

Is this the end-of-the-line for CPA as it stands?

Can it be run differently and live again?

Is it time to say that CSP is a good idea but no-one wants it?

Or is there some practical way forward that is worth the effort?

I will set-out, below, some thoughts on these questions. I do not necessarily have a view on how things should go but these notes are intended to provoke thought, by me and others.

Readers, please form your own views so that there can an exchange of ideas - sometime during the next CPA would be a good time.

I'm happy to turn this into a few slides as prompts for discussion - and also to add in other views if you let me have them (or you can just bring them to a discussion).

2 The origins of CPA

In the 1980's, a new computer was produced that combined computing, memory and communication on a single chip. It was possible to connect multiple chips as parallel processors and run complex algorithms at speeds unachieved by other, existing, systems.

At the time, this was world-leading and the chips made a significant impact. They were incorporated in many high-profile systems and are still, in a very small number of niche areas, being built into new products today (using stock parts, none have been made for many years).

Inmos, the company, and "Transputers", the chips, enabled many significant applications such as: graphics (e.g. ray-tracing in real-time), defence (radar tracking), instrument control, video-phones, ...

The highly successful devices were mostly 16- and 32-bit integer-only hardware (floating point in software libraries). A follow-on development of a device with floating-point hardware and faster communication links suffered political disruption and never made it to mainstream use.

The hardware efficiently implemented a new paradigm in the design of concurrent systems in which a system was built of a number of sequential processes that interacted via communications having well-defined semantics. Communicating Sequential Process paradigm leads to the CP of CPA.

Associated tools enable programs to be checked for good behaviour and reliable systems to be produced. A new programming language - occam - directly implemented the principles of CSP and, more importantly, allowed compilers to check important program characteristics directly.

In those heady-days, conferences were held in the UK, USA, Japan and Europe with audiences, briefly, up to 400 people.

3 Why did things go downhill?

Opinions are many and varied, here's my selection ...

Undoubtedly, the lack of Transputer hardware development was a major blow. While the original processor speeds up to 40MHz and the abortive successor reached 100MHz, traditional processors continued to advance to 4000MHz.

In becoming faster, sequential processors were able to overtake the performance of concurrent processor networks. The need for concurrent networks was limited to those who couldn't wait a year or two for sequential processor development to provide the needed performance.

occam was not seen to be acceptable to large parts of the programmer community. There has been a misguided desire to retain old sequential code rather than re-write it to exploit natural concurrency - there has probably been more effort expended in re-use than re-implementation. C has a stranglehold on programming - despite providing very little support for writing reliable code (how many of the headline security breaches have been caused by the simple error of read/writing off the end of an array?).

CSP seems to exist only in a small, and getting smaller, group of people (but see below for exceptions).

3.1 And What happened to CPA?

Perhaps in response to the apparent disappearance of CSP, the CPA call for papers no longer lists Communicating Processes as being of interest to the conference!

4 Is the Lack of Hardware Running CSP Directly Important?

It would certainly be nice to have and run very fast (remember what could be achieved by 40MHz processors and extrapolate to GHz machines) - but it is economically challenging (but not impossible) to develop new hardware.

On the other hand, existing silicon be used quite effectively (SPoC, KRoC, JCSP, CCSP, etc.) to build a market for new hardware - at least it could if there were some applications we could shout about - we need to identify some and shout about them.

There is an increasing number of combined processor / programmable logic devices available, from Xilinx and others, that could be used to add fast process-integrated communications to the CPU. The price/performance of a single device is not competitive, but how does a network stack-up? (I don't know, it needs some research.)

5 CSP: Signs of Death

- CPA papers are harder to get and are moving further from obvious use of CSP.
- CPA no longer asks for CP in submissions (not 'requires' but doesn't even suggest!).
- JCSP is not being (publicly) developed. From the very recently downloaded JCSP 1.1 notes: "There are rather a lot of deprecated items now in the library. Most of these will be removed in JCSP 1.2. This vew version will contain a completely re-written networking support (org.jcsp.lang.channel). Peter Welch and Neil Brown. (28th. November, 2008)"

  Six years with no change.
- RMoX hasn't made a splash.
- CSP-based web server hasn't had any noticeable impact.

6 CSP: Signs of Life

- CPA has presented some very interesting papers and continues to be attended by supportive people.
- Several (but not a huge number) Universities continue to do interesting things.
- Checking tools are still being developed (FDIR and others).
- CSP libraries are being ported to new languages.
- XMOS - successful in niche areas.
- Altreonic - a very small real-time kernel designed and proved with CSP - providing CSP to users.
- There is regular a meeting of people who want to use CSP to write better software - it takes place monthly and is attended by around 40 people. Our idea of a good annual conference takes place an order of magnitude more often - in Japan.
- Transputers are still being used. In some places it is old-stock and in others new development.
  - There is an FPGA Transputer in Japan.
  - There is a new chip developed in Europe.
- CSP networks are being built. 'SpaceWire' for on-board communications in satellites, although its usage is driven by non-CSP-aware designers leading to sub-optimal systems.

7 Why not just give up?

Reasons for slow take up of concurrency (development of ever-faster sequential processors) no longer apply. Concurrency **is** the way forward - if only it could be used properly. **NOW** is a good time to pass-on wisdom (although it may not easily be regarded as research, it is scholarship and can be academically acceptable).

There are now interesting ways to implement Transputers - as software, combined hardware/software (Xilinx ZynQ etc.) or even as new hardware (even relatively small companies can commission hardware - e.g. Parallella/Adapteva).

Interest in concurrency is growing - at both professional and hobbyist levels.

8 CSP: A way forward ...

I think we've rolled-on for some time now without a clear guiding purpose.

Now is a good time (actually, a few years ago was a good time - now we must catch-up) to consider where we are and where we are going (or would like to go).

- Review the current position
  - Where is work being done?
  - What is being done?
  - Where is it being applied (rumour as well as fact?). This could be the CSP Culture and/or actual, explicit, CSP. Note that it might be difficult to establish industrial use - if it is working and giving a competitive advantage they aren't going to admit it publicly (i.e. to their competitors)! We may not hear much from industry either because CSP is so bad it is no use or because it is so good it is being used profitably.
  - How is it being used? (Java, C, ...)
  - What is the state-of-play of JCSP/CCSP/SPoC/KRoC/etc.
  - Identify any successful applications of CSP (especially if they are significant examples that make a good case)
- Agree what benefits, if any, CP provides
  - Do we really have something to offer or are we just wanting others to agree with us to validate out biases?
  - Ease of transition from (a concurrent) real-world to a program
  - Proven behaviour in safety-critical systems (are we already too late to influence self-driving cars?)
  - ...
- Discuss how to disseminate 30-years work so the world can benefit from it:
  - Edit selected papers from WoTUG/CPA
  - Introduce CSP/occam/... into schools
  - Make CSP available on new devices (CPU, FPGA, ...) - new Transputers in some way
  - ...
- Where should interested parties meet?
  - At a conference (annual or otherwise)
  - At a workshop
  - Electronically - video conference, email, more interactive text exchanges, ...
- Who will do any of the identified work and how will the effort be justified / credited? (Sadly, it is all too easy to come up with a great plan but it has also to be a practical plan.)

8.1 Conferences

A traditional way to exchange ideas and meet people - with meeting being the most important bit.

Has CPA become just another conference on concurrency? If so, then it must be run in much the same way as others - meeting all the current measurement criteria (rejection rate, impact factor, etc. etc.) ... at least for UK academic contributors. Industry does not care much about these measures and I wonder if non-UK academics have the same measures to worry about. It must change considerably, in a way that most do not like but some have to live with.

If it is to be truly international (not just friends from the Netherlands and USA), it must be run to allow for political realities - e.g. be ready with public material (a list of accepted papers etc.) early enough for travel/attendance authorisation to be gained, visas can take weeks to months to obtain for those from e.g. Russia/China.

It is probably necessary to solicit contributions from the outset - not just at the submission deadline when things are looking thin.

8.2 Workshops

One very big advantage of a workshop is that papers can be sent to attendees in advance of the meeting (at least two weeks) so that further research can be done before the meeting and more detailed discussion can take place. (This is unlike conferences where papers are seen, at best, at the last minute.)

There is a theme, discussion is focussed and fewer but longer presentations/demonstrations are possible.

Since the early days, there have been few workshops - but one that should be remembered (the last one?) took place at UKC many years ago. It is notable that, as a result of the discussion on how to run occam on machines other than Transputers, TWO occam compilers were produced (SPoC and KRoC). These have seen extensive use, including industrially, and KRoC has been used as a basis for considerable research (I believe SPoC also produced some significant research output, but I'm less familiar with its use).

Because the last one was so successful, I'd like to see serious consideration of workshops.

Could a report be published (with suitable ISBN)?

The big question here is whether academics could attend - what is the position with regard to funding etc.?

8.3 Electronic meetings

These are relatively low-cost options.

My experience of video conferencing and email/other exchange has not been very good - even with as few as three people (real-meeting cues are absent, taking turns doesn't happen as you'd expect, time-delays in the communication make for very stilted conversations, ...).

A big problem is that they can become too open and too wide a range of views becomes distracting.

9 Final Remarks

In many ways, things are depressing - it is very unlikely we can continue in the same way, even for one more time.

However, there are great opportunities to do something worthwhile - if only we can work what to do and find some way to do it.

Let's analyse the situation and either stop for good reason or carry on with a realistic plan that leads to results we can be pleased to be part of.