Reply to Referees

Martin Berger

First I would like to thank all involved for their thorough work. One of the referees found a marvelous counterexample to my characterisation theorem and I'm very grateful for this. It shows that the timed π-calculus is even more interesting than I though. The counterexample affects only the completeness claim. All the proposed labelled equivalences are still sound. Furthermore, the counterexample and all additional ones I have come up with are pathological in the sense that one would not encounter them in programming and verification practice. As the paper will have to be rewritten from scratch to deal with the failure of completeness, I will only comment on the counterexample and on how to solve the problem it brings in focus. Please recall that I introduced an asynchronous transition system and claimed that the induced largest bisimulation that is closed under time-stepping and arbitrary renaming coincides with the unlabelled reduction congruence. Unfortunately this is not the case. There are reduction congruent processes that aren't bisimilar with respect to all the transition systems I introduced. I shall now present a simplified variant of this counterexample here, which has fewer transitions and reductions, and may hence be easier to understand.

The Calculus. For reference I start with some key definitions of the timed asynchronous π-calculus. Processes are given by the following grammar.

$$P ::= 0 \ \mid\ \overline{x}\langle\vec{y}\rangle \ \mid\ x(\vec{v}).P \ \mid !x(\vec{v}).P \ \mid\ \tau^t(x(\vec{v}).P, Q) \ \mid\ P|Q \ \mid\ (\nu x)P$$

The structural congruence has the usual definition, but on a larger set of processes. Reduction semantics is given next.

$$\textsc{Com}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|x(\vec{v}).P \to P[\vec{y}/\vec{v}]} \qquad \textsc{Rep}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|!x(\vec{v}).P \to P[\vec{y}/\vec{v}]|!x(\vec{v}).P}$$

$$\textsc{TCom}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|\tau^t(x(\vec{v}).P, Q) \to P[\vec{y}/\vec{v}]} \qquad \textsc{Par}\ \frac{P \to Q}{P|R \to Q|\phi(R)} \qquad \textsc{Res}\ \frac{P \to Q}{(\nu x)P \to (\nu x)Q}$$

$$\textsc{Idle}\ \frac{-}{P \to \phi(P)} \qquad \textsc{Cong}\ \frac{P \equiv P' \to Q' \equiv Q}{P \to Q}$$

The *time-stepper* function ϕ has this definition.

$$\begin{array}{rcl@{\qquad}rcl}
\phi(\tau^1(x(\vec{v}).P, Q)) & = & Q & \phi(\tau^{t+1}(x(\vec{v}).P, Q)) & = & \phi(\tau^t(x(\vec{v}).P, Q)) \\
\phi(P|Q) & = & \phi(P)|\phi(Q), & \phi((\nu x)P) & = & (\nu x)\phi(P) \\
\phi(P) & = & P \text{ for all other } P. & & &
\end{array}$$

The relation $\simeq^{rc}$ is the largest congruence preserving asynchronous barbs and reduction steps. It is irrelevant to the counterexample if one take all contexts or just reduction contexts, but for simplicity I choose the former here.

Transitions Next I present a *synchronous* transition system that is slightly different from those presented in the TCS submission and my thesis. I developed it in response to the referees' comments and think it's neater than its predecessors. However, the counterexample below will also be virulent for all other proposed transition systems. The idea behind this transition system is that it's exactly like the usual synchronised one, but having a dedicated action time-passing action ϕ which can also match τ, but not vice versa, in bisimulations.

The *labels* are generated by the grammar

$$l ::= \tau \mid \phi \mid x(\vec{y}) \mid \overline{x}\langle(\nu\vec{y})\vec{z}\rangle$$

The relation $\unrhd$ is the least reflexive relation on labels such that $\tau \unrhd \phi$. Transitions $\xrightarrow{l}$ are inductively generated by the rules in Figure 1. Then one defines $P \overset{l}{\rightsquigarrow} Q$ iff

$$P \xrightarrow{l} Q, \text{ or } \qquad l = x(\vec{y}), P \xrightarrow{\phi} Q', P \equiv Q'|\overline{x}\langle\vec{y}\rangle.$$

The relation $\unrhd$ and the auxiliary transition $\overset{l}{\rightsquigarrow}$ help to ensure that 0 and forwarders fw_x are equated by the induced labelled bisimilarity. This is of course crucial for completeness, but irrelevant for the counterexample. Now let $\sim$ be the largest relation such that $P \sim Q$ implies (1) $P\sigma \sim Q\sigma$ for any renaming σ and (2) if $P \xrightarrow{l} P'$ then there is $Q \overset{l'}{\rightsquigarrow} Q'$ such that $l \unrhd l'$ and $P' \sim Q'$, and vice versa.

The Counterexample. I begin by defining some auxiliary processes that will prove useful later.

$$\mathsf{delay}^0(P) \overset{def}{=} P \qquad \mathsf{delay}^{t+1}(P) \overset{def}{=} (\nu x)(\overline{x}|\tau^t(y.0, P)) \qquad (x \text{ fresh}).$$

Then $\mathsf{delay}^{t+1}(P)$ has exactly one transition.

$$\mathsf{delay}^{t+1}(P) \xrightarrow{\phi} \mathsf{delay}^t(P).$$

For a finite collection $(P_i)_{i\in I}$ the *timed sum* is

$$\bigoplus_{i\in I} P_i \rhd Q \overset{def}{=} (\nu x)(\overline{x}|\Pi_{i\in I}\tau^1(x.P_i, Q)) \qquad (x \text{ fresh}).$$

A special case of this last definition is

$$P \rhd Q \overset{def}{=} (\nu x)(\overline{x}|\tau^1(x.P, Q)) \qquad (x \text{ fresh}).$$

$$\textsc{Out}\ \frac{-}{\overline{x}\langle\vec{y}\rangle \xrightarrow{\overline{x}\langle\vec{y}\rangle} 0} \qquad \textsc{In}\ \frac{-}{x(\vec{v}).P \xrightarrow{x(\vec{y})} P[\vec{y}/\vec{v}]} \qquad \textsc{RIn}\ \frac{-}{!x(\vec{v}).P \xrightarrow{x(\vec{y})} P[\vec{y}/\vec{v}]|!x(\vec{v}).P}$$

$$\textsc{TimeIn}\ \frac{-}{\tau^t(x(\vec{v}).P, Q) \xrightarrow{x(\vec{y})} P[\vec{y}/\vec{v}]} \qquad \textsc{Com}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|x(\vec{v}).P \xrightarrow{\tau} P[\vec{y}/\vec{v}]}$$

$$\textsc{Rep}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|!x(\vec{v}).P \xrightarrow{\tau} P[\vec{y}/\vec{v}]|!x(\vec{v}).P} \qquad \textsc{TCom}\ \frac{-}{\overline{x}\langle\vec{y}\rangle|\tau^t(x(\vec{v}).P, Q) \xrightarrow{\tau} P[\vec{y}/\vec{v}]}$$

$$\textsc{Par}\ \frac{P \xrightarrow{l} Q \quad \mathsf{bn}(l) \cap \mathsf{fn}(R)}{P|R \xrightarrow{l} Q|\phi(R)} \qquad \textsc{Res}\ \frac{P \xrightarrow{l} Q \quad x \notin \mathsf{n}(l)}{(\nu x)P \xrightarrow{l} (\nu x)Q} \qquad \textsc{Idle}\ \frac{-}{P \xrightarrow{\phi} \phi(P)}$$

$$\textsc{Open}\ \frac{P \xrightarrow{\overline{a}\langle(\nu\vec{b})\vec{c}\rangle} Q \quad x \in \vec{c} \setminus \{a, \vec{b}\}}{(\nu x)P \xrightarrow{\overline{a}\langle(\nu\vec{b}x)\vec{c}\rangle} Q} \qquad \textsc{Cong}\ \frac{P \equiv P' \xrightarrow{l} Q' \equiv Q}{P \xrightarrow{l} Q}$$

Figure 1: The synchronous transition system.

For $I \neq \emptyset$, the timed sum has exactly one transition more than I has elements.

$$\bigoplus_{i \in I} P_i \triangleright Q \xrightarrow{\phi} Q \qquad \bigoplus_{i \in I} P_i \triangleright Q \xrightarrow{\tau} P_i.$$

I shall also write $\overline{x}(\vec{y})P$ for $(\nu\vec{y})(\overline{x}\langle\vec{y}\rangle|P)$, provided $x \notin \vec{y}$.

Next are two processes P and Q which are contextually equivalent but nevertheless distinguished by all the obvious labelled equalities, assuming that x, y, a are distinct names.

$$P \overset{def}{=} \overline{x}(yr)r. \bigoplus_{i=1,2,3} P_i \triangleright 0 \qquad\qquad Q \overset{def}{=} \overline{x}(yr)r. \bigoplus_{i=2,3} P_i \triangleright 0$$

The component processes P_i are given next.

$$P_1 \overset{def}{=} y.(\overline{a}|\overline{y}) \qquad P_2 \overset{def}{=} \overline{a} \triangleright P_1 \qquad P_3 \overset{def}{=} \mathsf{delay}^1(P_1).$$

It is easy to see that $P \not\approx Q$ because P_1 can output $\overline{a}$ only after inputting y, behaviour that can neither be matched by P_2 nor by P_3. But why would $P \simeq^{rc} Q$ hold? The key question to answer here is how to match

$$C[(\nu y)(R| \bigoplus_{i=1,2,3} P_i \triangleright 0)] \rightarrow \phi(C)[(\nu y)(\phi(R)|P_1]. \tag{1}$$

Clearly

$$C[(\nu y)(R| \bigoplus_{i=2,3} P_i \triangleright 0)] \rightarrow \phi(C)[(\nu y)(\phi(R)|P_j]$$

is inappropriate for both, $j = 2$ and $j = 3$. The key insight about why $P \simeq^{rc} Q$ nevertheless holds is that

$$C[(\nu y)(R|P_1)] \quad \simeq^{rc} \quad C[(\nu y)(R|P_2)] \tag{2}$$

for all $C[\cdot]$ and R, provided R does not have a strong barb at y. If on the other hand R does have such a strong barb, then

$$C[(\nu y)(R|P_1)] \quad \simeq^{rc} \quad C[(\nu y)(R|P_3)] \tag{3}$$

for all $C[\cdot]$ and such R. But in the process on the left of (1) one already knows if $\phi(R)$ will have a strong barb at y or not. It is already decided for each given $C[\cdot]$ and R. Hence reduction (1) can be matched by either (2) or (3), depending on $C[\cdot]$ and R.

Of course R can decide if it wants do provide an output or not. But in the present calculus, that decision *takes time*, at least one unit and the timed sums in P and Q are such that the difference (that makes them non-equal by simple labelled bisimilarities) between the two is very transient: it disappears after one unit of time, because of judiciously set timers. The reason similar phenomena do not occur in untimed calculi like the asynchronous π-calculus is that making an internal choice does not immediately affect the environment there.

Proposed Research. Is it possible to overcome this problem? My answer at this point is a cautions "possibly". I have a notion of labelled bisimilarity that does equate P and Q. But it is relatively straightforward to construct more complicated counterexamples, for example by iterating the trick embodied in P and Q, that would exhibit the original problem.

Let's look at that new definition. Assume $\subseteq_f$ stands for finite subsets and $\mathsf{ex}(l)$ denotes the set of names that are output bound by l. A collection $(\mathcal{R}_S)_{S \subseteq_f \mathcal{N}}$ is a *bisimulation* if each $\mathcal{R}_S$ is an S-*bisimulation*, that is whenever $P \,\mathcal{R}_S\, Q$ then (1) $P\sigma \,\mathcal{R}_{\sigma[S]}\, Q\sigma$, and (2) whenever $P \xrightarrow{l} P'$ and $\mathsf{ex}(l) = \emptyset$, then at least one of the following is the true.

- There is $Q \overset{l'}{\leadsto} Q'$ such that $l \trianglerighteq l'$ and $P' \,\mathcal{R}_{S \cup \mathsf{ex}(l)}\, Q'$.
- If $l = \tau$ and P' has exactly two types of transitions:
  - $P' \xrightarrow{\phi} P_\phi$ and
  - for all $x \in S$ and all $\vec{a}$: $P' \xrightarrow{x(\vec{a})} P_{x(\vec{a})}$.

  For all those it is possible to find

  $$Q \overset{l'}{\leadsto}\xrightarrow{\phi} Q_\phi \qquad P_\phi \,\mathcal{R}_S\, Q_\phi$$

  and

  $$Q \overset{l'}{\leadsto}\xrightarrow{x(\vec{a})} Q_{x(\vec{a})} \qquad P_{x(\vec{a})} \,\mathcal{R}_S\, Q_{x(\vec{a})}.$$

Now $\sim \overset{def}{=} (\sim_S)_{S \subseteq_f \mathcal{N}}$ denotes the largest bisimulation, computed pointwise.

One can easily show that $P \sim_{\emptyset} Q$ and $\sim_{\emptyset} \subset \simeq^{rc}$. The definition just proposed seems somewhat arbitrary in that it allows to weaken the strict matching discipline of conventional bisimilarities a bit, by considering traces of length two, under certain restricted conditions. Why length two only? I do believe, and have made progress towards proofs that $(\sim_S)$ can be generalised quite nicely and deal with a much larger class of counterexamples. However, whether such a generalisation would lead to completeness is unclear. There are four questions to which I currently have only partial answers.

1. In the example, and all it's variants, the unmatched transition is an input on a name that was extruded boundly earlier. Are similar phenomena possible with a free name z? I rather don't think so, because one can always place a critical process in a context $[\cdot]|\overline{y}$, so it is never guaranteed that z is not available immediately.

2. A similar question may be asked with respect to inputs vs. other types of actions. Does the action not matched have to be an input? I don't know, but suspect that it may have to.

3. All the examples also use the critical name in subject position. Is it possible to come up with a counterexample where this name is only an object? My guess is that this should not be possible, because when one uses a name, it may be supplied by the context.

4. All examples I could come up with are derived from the process

   $$x(\vec{v}).(\overline{x}\langle\vec{v}\rangle|P).$$

   It is a generalisation of a forwarder in that it does not visibly consume its initial input. Might unrelated processes that do not have this peculiar property lead to similar mismatches between reductions and simple-minded transitions? Here I'd venture a two part guess.

   - Forwarder-like processes are unique in that they are equated to processes that have different visible initial actions in synchronous transition systems. There are no other such processes.
   - Each instance of a mismatch between labelled and unlabelled semantics arises from a mismatching initial visible action like with forwarders.

5. In all the counterexamples the trace-length of the behaviour that cannot be matched conventionally is finite, in the original example it has length 2. But would that be necessarily the case or is it possible to have processes with infinitly long traces that cannot be matched, and yet these processes are contextually indistinguishable? I suspect it must be possible, because if P has an infinite ϕ-free trace, the outputs available for that trace are already decided upon when P becomes active, hence we should

> be able to match with different transitions for different decisions even in the infinite case. It's just that such an infinite process would be complicated, especially when using replication rather than recursive equations[1]. Another difficulty is that it becomes unclear how to avoid degeneration into a trace-like equivalence if traces no longer end with future states that must be matched. That problem might be solvable with approximations and limits, but I hope simpler solutions can be found.

Conclusion. Soundness and usefulness of the proposed labelled notions of equivalence is not affected by the unearthed counterexamples. Nevertheless it is an interesting technical challenge to look for (1) bisimilarities that can equate all the counterexamples suggested by the referee's initial one, and (2) to consider completeness. I believe that the former is feasible with some concerted effort. The new notion of bisimilarity sketched above seems to have a neat generalisation that can equate all the counterexamples that are like the referee's. I will report on this soon. Whether this is enough for completeness is unclear. I am hopeful but there may be radically different classes of counterexamples. I'm currently investigating the matter.

[1] This touches on another unresolved issue: how to connect timing and replication/recursion in a satisfactory manner. The solution chosen in my thesis and the CONCUR sumission, constant time replication, is but a convenient starting point.