The Ontology of States, Processes, and Events

Antony Galton

College of Engineering, Mathematics and Physical Science, University of Exeter, UK
apgalton@ex.ac.uk

ABSTRACT. This paper presents a new view of the relationship between states, processes and events. Instead of trying to treat them as entities all on a similar footing, as most previous authors have done, we regard processes as abstract patterns of behaviour which may be realised in concrete form as actually occurring states or events. Processes are divided into two broad types, called *continuables* and *repeatables*, and various mappings between and within these categories are considered. The theory presented here is consistent with recent theorising about processes in ontology and computer science while being sensitive to the insights from the work of philosophers and linguists over many years.

1 Introduction

According to [11], "processes are repeatable behaviours whose occurrences cause continuants to undergo change". In this paper I present a view of processes which takes as its starting point part of this idea — that is, the notion of processes as "repeatable behaviours" which can have "occurrences" — and develops it in a way that is sensitive to the insights arising from the study of aspect and aspectual character which has for many years formed an important part of theorising about states, process, and events.

The major problem, as I see it, has been to fit the notion of *process* into an account of temporal phenomena which is ruled by an overarching dichotomy between *states* and *events*. The difficulty that this task presents is well illustrated by the existence of several incompatible attempts by researchers from various disciplines to come up with a satisfactory taxonomy of the things that go on in time [3]. Some examples are shown in diagrammatic form in Figure 1.

It will be seen from these examples that there is little agreement about how states and events are related to each other, and even less about how processes are related to states and events. This diversity arises at least in part from terminological differences: the authors are simply using these terms as labels for different things. For example, comparing Allen [1] and Pustejovsky [12], it seems likely that Allen's "event" is similar to Pustejovsky's "transition", whereas Pustejovsky's "event" is a broader term, with no equivalent in Allen's terminology, covering all three of Allen's categories. But the disagreements can scarcely be purely terminological, otherwise we might at least expect to see greater uniformity in the structures of the taxonomies presented.

In this paper, I advocate a point of view which is in some respects utterly at variance from those expressed by the authors represented in Figure 1 — and indeed from my own

previous publications on this subject. According to this new point of view, processes are a radically different *kind* of entity from states and events, so any attempt to bring the three categories together in a simple subsumption hierarchy such as those illustrated in Figure 1 is doomed to failure. In a nutshell, processes are abstract entities which may be *realised* as concrete entities which are actually occurring states and events. The precise ontological status of such abstract entities is problematic to the extent that abstract entities in general may be considered to be so, but I believe that, however this particular dilemma is played out, the relationship between processes, states, and events proposed in this paper will remain viable.

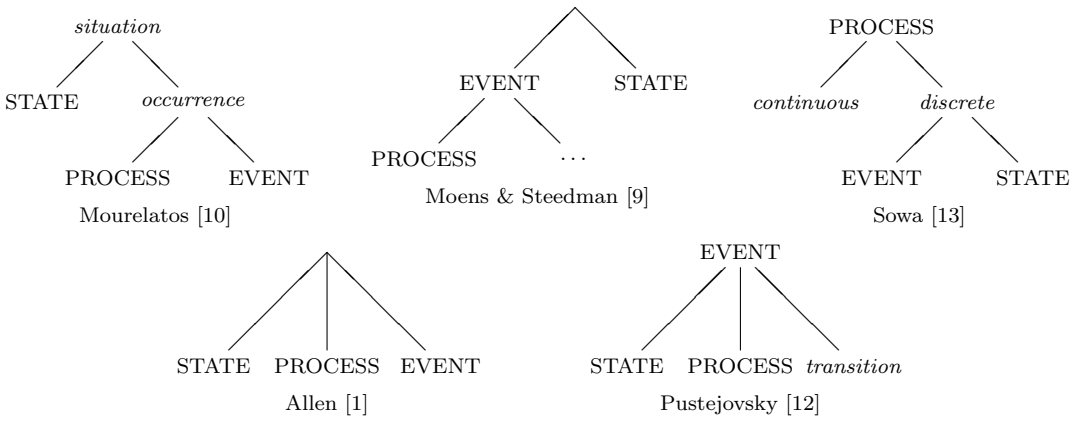

Fig. 1. Some taxonomies of temporal phenomena

2 Types and tokens

The literature on events makes extensive use of the distinction between *event types* and *event tokens*. Event tokens are supposed to be actual chunks of spatio-temporal reality, involving real things undergoing real changes or movements. An event token has a more or less definite location in time and, subject to various caveats about the location of events with spatially distributed participants, in space. Event types are supposed to be "abstract", or "universal", entities, of which event tokens are *realisations* or *instances*. An event token is often said to be an *occurrence* of an event type: thus the terms "event token" and "occurrence" may often be interchangeable.

People sometimes try applying this type/token distinction to processes and states as well as events. This is much less straightforward. Consider a simple more or less homogeneous, open-ended process such as *walking*. This seems to have a "universal" character to it in the sense that we can certainly find walking at many different spatio-temporal locations and with many different participants. For example, at 10.40 a.m., on Sunday 1st January 2012, I was walking near my house in Exeter. It is tempting to think that there is therefore a walking *token* located at that time. But if so, what is the

spatio-temporal extent of that walking? The walking I was doing at 10.40 a.m. was a continuation of the walking I had been doing over the preceding few minutes; and was itself continued for a few minutes more. In fact I was walking from my house to the local railway station. It seems natural to pick out the entire walk, starting from when I left my house and finishing when I arrived at the station, and call this a walking token. But as soon as we have defined it in this way, what we have is surely an *event* token, a *walk* by me on a particular occasion. This can be described in various different ways, depending on which event type it is assigned to. In ordinary discourse we typically report the occurrence of such an event using a sentence in the perfective aspect, for example "I walked to the station", or "I walked for five minutes", or "I walked 500 metres". It is not clear that we need a notion of "process token" that is distinct from the widely accepted notion of event token.

On the other hand we might focus attention on the walking that I was doing at precisely 10.40 a.m. For it is true that I was walking then: a snapshot of the relevant part of Exeter at that time would include a representation of a man walking. Maybe it is this, the walking by me at 10.40, that should be labelled as a process token? But that I was walking at that time seems to amount to no more than that I was in a particular *state* at that time, that is, that the various parts of my body were all in particular physical and physiological states (including states of motion or change) that are characteristic of a particular phase of walking. If this is a token of anything, it would appear to be sensible to describe it as a state token. That such a token exists at a particular time is typically reported using a sentence in the imperfective aspect, for example "I was walking", or, with a particular "colouring" obtained by describing the state in relation to an event type which it contributes to the realisation of, "I was walking to the station", "I was walking for five minutes", or "I was walking 500 metres".

This leads us to a picture in which we have an event token — my complete walk to the station — which occupies a certain five-minute interval of time, and an indefinite number of state tokens — all the momentary states of walking — that in aggregate give rise to the complete walk.

Where is the process in all this? Following the cue of [11], we might say that walking is a "repeatable behaviour" of which the actual state and event tokens we find in the world are "occurrences". More exactly, we should distinguish between *repeatable* behaviours such as "walking to the station", "walking for five minutes" or "walking 500 metres", whose occurrences are actual events (i.e., event tokens), and *continuable* behaviours such as "walking" pure and simple, which can be realised over arbitrarily short time intervals and indeed can meaningfully be ascribed to arbitrary moments within any interval over which it is realised — and in that case what it is instantiated as the state that obtains then by virtue of the particular sequence of movements and/or changes of which it forms a part.

This discussion has glossed over many of the niceties that are required for a more adequate account. It represents a first foray into the territory which will be delineated in greater detail in the remainder of this paper.

3 Continuables and Repeatables

At the end of the last section I drew a distinction between behaviours that are *repeatable*, i.e., realisable in the form of discrete occurrences at different times, and behaviours that are *continuable*, i.e., realisable in the form of a momentary state that can evolve continuously over an indefinite period of time. In this section I will build up a more elaborate system of distinctions taking this as the starting point.

We start with the idea of a *simple generic continuable*. This is a homogeneous, open-ended behaviour which may be *enacted* by an agent or set of agents over a period of time. The terms 'enact' and 'agent' are not to be understood here as implying intentional agency: the agents can be physical entities responding to imposed forces, as well as intentional agents capable of initiating actions autonomously [4]. Homogeneity is, as always, to be understood with respect to some chosen level of granularity. Simple generic continuables are what are denoted by simple verbs such as "run", "sing", "eat", or "whistle". They are abstract in the sense that they do not in themselves exist in space or time or have any kind of material reality; rather, they are patterns which may be realised at different times and places.

Simple generic continuables can be made more specific, while still retaining their continuable character, in two main ways. The first possibility is to specify a *non-delimiting* object, that is, a noun or noun phrase which designates its referent in an open-ended way. Non-delimiting objects are typically either mass terms or indefinite plurals, e.g., "apple" or "apples", as opposed to "an apple", "this apple", or "five apples". This kind of continuable, then, is denoted by a verb phrase consisting of a simple verb together with a non-delimiting object, e.g., "sing songs", "sing Schubert" (here "Schubert" functions as a mass noun!), "eat apple", "eat apples". In these examples, the verb is transitive so it is possible to specify an object. However, other forms of non-delimiting qualification can be applied even in cases where the verb is intransitive, e.g., "run northwards".

The other way of making a simple generic continuable more specific is by specifying an agent or set of agents for it. A *simple specific continuable* is derived from a simple generic continuable by restricting it to a particular agent or set of agents. Such a continuable is still abstract, i.e., not located in space and time, but in one sense it is already more concrete since the agent is. Simple specific continuables are denoted by verb phrases consisting of a simple verb together with a subject and, optionally, a non-delimiting object as discussed in the previous paragraph. Examples are: "John runs", "Mary sings", "Mary sings songs", "Mary sings Schubert", "Bill eats", "Bill eats apple(s)", "The kettle whistles".

We can combine simple continuables to produce *compound continuables*. For reasons that will become clearer later, the only form of combination allowed for continuables is *simultaneous composition*, under which two continuables are combined to form a continuable whose realisations are precisely the simultaneous realisations of the two constituent continuables: the general form for denoting this may be given as "X while Y", where X and Y are the continuables to be combined. If X is generic, so is the compound, even if Y is specific: thus generic compound continuables include examples such as "read while eating", "read while John eats", and "read poetry while John eats apples". If X is specific, then Y must also be specific, and the compound is then specific too, e.g., "John reads while Mary eats". We shall describe composition as *homogeneous*

when it is applied to two elements of the same specificity (i.e., either both generic or both specific), and *heterogeneous* when applied to elements differing in specificity.

We thus have four broad classes of continuables, related as shown in Figure 2.

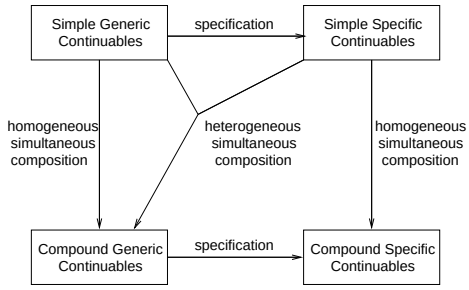

Fig. 2. Four classes of continuables

The next stage in the development is to introduce the class of *repeatables*. As noted above, a continuable can be made more specific by the addition of a non-delimiting qualification such as a direct object or a direction. The result in such cases is still a continuable. If, however, we attach a *delimiting* qualification to a continuable, the result is no longer a continuable. This is because the delimiting qualification introduces the idea of *completion*, and something that has been completed is no longer continuable: if I finish eating an apple, I cannot continue eating it. On the other hand, I can then eat another apple, and this shows how the notion of repeatability comes in here. The result of adding a delimiting qualification to a simple continuable is therefore a *simple repeatable*. This will be generic or specific according as the continuable from which it is derived is. Examples of simple generic repeatables are "run to the station", "run a mile", "sing for an hour", "sing *Happy Birthday*", and "eat an apple". Note that the delimiting qualification can be a direct object or an adverbial phrase; the important point is that it specifies what counts as a completion for the repeatable in question. Examples of simple specific repeatables are "John runs a mile", "Mary sings for an hour", "Mary sings *Happy Birthday*", and "Bill eats an apple". Note that these specific repeatables can already be described as designating *event types*.[1]

Delimitation can, in principle, be applied to compound continuables, resulting in compound repeatables, but this operation needs to be handled with care. A straightforward case is when the delimitation is of a kind that can be applied to arbitrary continuables, e.g., "for an hour". Thus from the compound continuable "John bakes while Mary writes" we can form the compound repeatable "For an hour, John bakes while Mary writes". In other cases, we might want to apply different delimiting qualifications to the two parts of a compound continuable, e.g., "John bakes a cake while

[1] Not all repeatables can in fact occur more than once, e.g., "assassinate Caesar". This is not a flaw in the theory so much as a less than ideal terminology; another possibility would be to use the term "completable" for what is here designated as "repeatable", although this also may be less than ideal.

Mary writes an essay". It is, however, more natural to think of this as obtained by forming a compound of the two repeatables "John bakes a cake" and "Mary writes an essay".

Repeatables admit of a wider range of possible modes of composition than just the simultaneous composition that is applicable to continuables. In the first instance, one might think of *sequential* and *parallel* composition. In the former, one repeatable is enacted first, and the other one after the first one is finished. With parallel composition, they are enacted at the same time. But "at the same time" is rather ambiguous when applied to occurrences which may be of different durations; instead we should consider explicitly the possible ways in which two occurrences can overlap. For this purpose we have a ready-made theory in the form of the Interval Calculus [1].

Allen [1] enumerated the 13 qualitatively different ways in which one interval can be temporally related to another. One of the 13 is strict equality ("A equals B"). The other 12 form six pairs of mutually inverse relations. The illustration in Figure 3 shows one relation from each pair, together with equality; the symbols used to denote the relations are shown along the bottom row of the diagram; for the inverse relations, we write, e.g., "si" for "is started by". The Interval Algebra includes, in addition to these 13 relations, all possible sets of such relations, interpreted as disjunctions. Thus, with "m" interpreted as "meets" and "o" interpreted as "overlaps" (in the specific sense illustrated in the figure), the relation "{m,o}" can be read as "meets or overlaps".

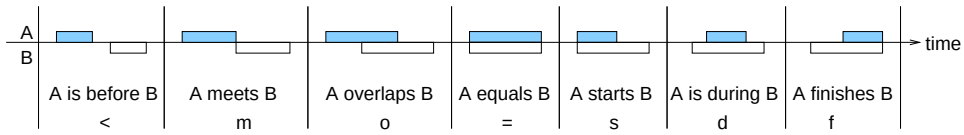

Fig. 3. Seven of the 13 relations of the Interval Calculus

Composition of repeatables can now be specified in the form XRY, where X and Y are the two repeatables to be composed, and R is one of the relations from the Interval Algebra. Thus we can write

$$\text{(John bakes a cake)}\{=,\text{s},\text{si}\}\text{(Mary writes an essay)},$$

which denotes the compound repeatable each of whose occurrences consists of an occurrence of John baking a cake and an occurrence of Mary writing an essay, where the former occurs over an interval which equals, starts, or is started by the interval over which the latter occurs — in short, the two occurrences begin simultaneously, it being unspecified which, if either, ends first. We will use the term *interval composition* to apply to this kind of composition. As with simultaneous composition of continuables, interval composition of repeatables can be homogeneous or heterogeneous depending on the specificities of the elements being composed.

With this apparatus in place, we can often avoid having to apply delimitation to compound continuables, since the same effect can be obtained by composing repeatables resulting from delimiting the constituent simple continuables. Our earlier example, "For

an hour, John bakes while Mary writes", which we obtained by first composing the simple continuables "John bakes" and "Mary writes" and then applying the delimiter "for an hour" to the compound, can equally well be obtained by applying the delimiter to the simples, and then composing them using "equals", giving:

(John bakes for an hour){=}(Mary writes for an hour)

However, this reduction does not apply in all cases. Consider "John bakes a cake while singing". This is a repeatable, not a continuable, since the delimitation applied to "bake" overrides the non-delimitedness of "sing". But it cannot be obtained by composing two repeatables. One possibility might be to allow composition of a repeatable with a continuable; but we can avoid this by applying delimitation to one component of a compound continuable but not the other: thus "John bakes a cake while singing" can be regarded as the result of applying the delimiter "a cake" to the first component of "John bakes while singing".

We are now in a position to extend our diagram to include repeatables. The result is illustrated in Figure 4.

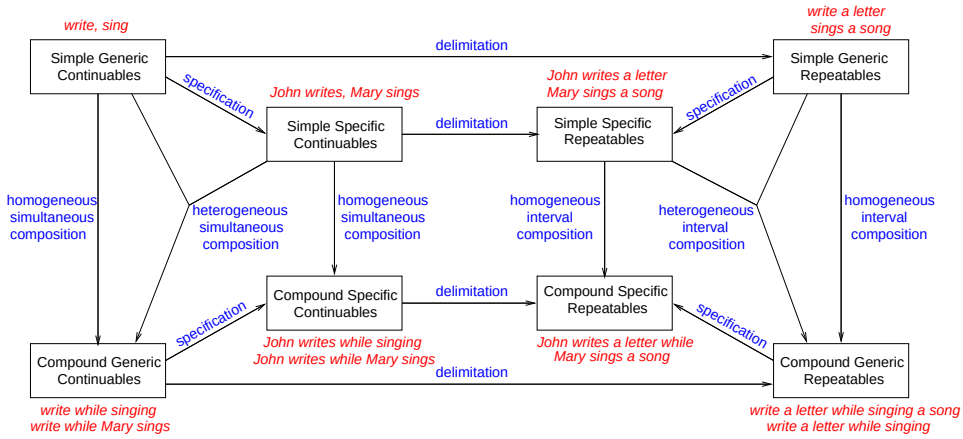

Fig. 4. Continuables and repeatables

4 Realisations

So far, everything we have considered has been at the abstract level: we have been dealing with types of activity or event that can be instantiated on different occasions and in different ways. Such instantiations are concrete realisations of these continuable or repeatable behaviours. As such they are fully determinate with respect to their spatial, temporal, and indeed all other characteristics. The continuables and repeatables discussed in the previous section provide a repertoire of descriptive resources that can

be used to specify the form of actual states and events. Here the terms "event" and "state" are used to refer to what actually happens or is the case; they correspond to two fundamentally different perspectives we can take on the world [3].

In the case of states, we are taking the perspective of someone immersed in the world, participating in it at a time: the state of the world comprises everything that is captured in a "snapshot" of the world at a time. It must be emphasised that this includes the dynamic characteristics of the world: if something is moving or changing at a time t, then this fact is included in the state of the world at t. For this reason the term "snapshot" is misleading since it suggests that, as with a camera, all motion is "frozen". But the world does not consist of a succession of states in which everything is motionless! [3, 4]

In the case of events, what happens is being viewed from a perspective which spans a period of time. It is the synoptic view one obtains by projecting oneself outside the immediate here and now and taking a historical perspective. We are able to take such a view because of human memory and external recording media. These are essential, since no event which takes more than a very short time is ever present to a person's immediate experience in its entirety.

The use of the terms "snapshot" and "spans" in the previous two paragraphs recalls the SNAP and SPAN sub-ontologies of BFO [5, 6]. The inhabitants of the SNAP domain are *continuants*, those entities which can exist at individual times and endure over periods of time; the BFO ontology includes states as SNAP entities, although it is not entirely clear whether these states can be dynamic as required by the theory presented here. The SPAN domain comprises entities which cannot completely exist at a time but rather by nature span periods; these are *occurrents*. Events are included as SPAN entities in BFO: but the name given to them there is "processes".

Abstract continuables and repeatables are mapped onto states and events by the operation of *realisation*. Continuables are realised as states and repeatables as events. If I say "Mary is singing", I mean that the current state of the world includes a realisation of the simple specific continuable "Mary sings". If I say "Mary sang *Happy Birthday* yesterday", I mean that that part of the history of the world which I designate "yesterday" includes a realisation of the simple specific repeatable "Mary sings *Happy Birthday*". The top three rows of the diagram in Figure 5 show the various mappings involved in building up these descriptions from basic ingredients and relating it to actual circumstances in the world which realise them.

On the picture painted thus far, continuables are always realised as continuants (states), and repeatables as occurrents (events). But there are cases where it might seem that a continuable should be realisable as an occurrent, and a repeatable as a continuant. As an example of the former type, suppose I say "Mary sang twice yesterday". The adverb "twice" clearly indicates that what is in question is a repeatable with two realisations: there are two episodes of singing. But "Mary sings", by our account, is a continuable. In order to handle this, we must introduce a "null delimiter", which converts a continuable to a repeatable which denotes an *episode* of the activity denoted by the continuable. An episode of Mary's singing is a period in which Mary begins to sing, then sings for a while, then stops singing. In [2] I used the operator Po to convert a state into the event whose occurrences are episodes of that state. Thus Po acts as a form of delimiter and can be used to map the continuable "Mary sings" onto

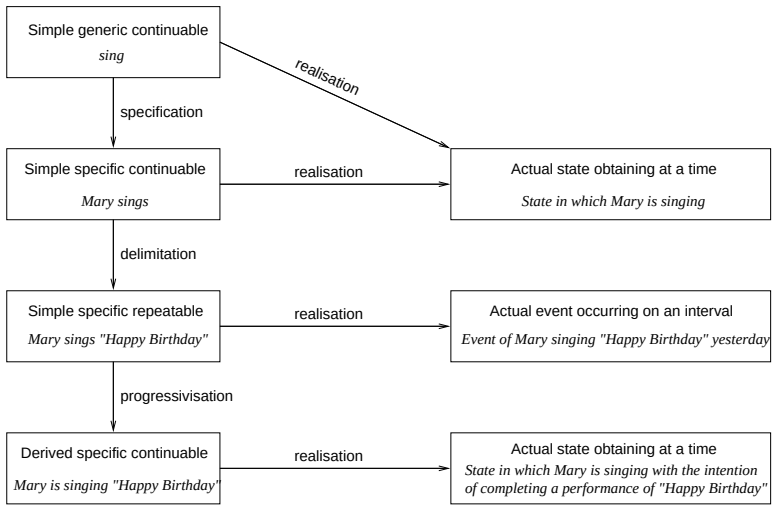

Fig. 5. Mappings relating continuables, repeatables, and their realisations

a repeatable "Mary sings for a while", where the qualifier "for a while" is intended to be as general as possible, covering any episode of Mary's singing without reference to how long or short it is, or any other characteristics of it. By using the delimiter *Po*, we dispel the appearance that a continuable can be realised as an occurrent: in a sense it can, but only by first being converted to a repeatable.

The other kind of anomaly, a repeatable realised as a continuant, is illustrated by sentences in the progressive aspect such as "John is baking a cake", which clearly denotes a state rather than an event. But it is a state that is intimately related to the event denoted by the repeatable "John bakes a cake". An important difference is that whereas the event involves a completion — there is a cake at the end of it — the state denoted by the progressive form does not imply completion. The solution is again to invoke an operator, this time mapping a repeatable to a continuable which holds by virtue of the repeatable being *in progress*. What "in progress" means has to be spelt out in terms of the means by which the continuable can be regarded as *progressing* or *directed towards* the completion of the repeatable. This is discussed in detail in [2], in the analysis of the progressive aspect. There are three main ways in which such directedness can arise:

1. *Intention*. The continuable is an activity of an intentional agent who *intends* that the activity should be (part of) what brings about the completion of the repeatable.
2. *Causation*. The continuable is such that, under the normal operation of causal laws, it will be (part of) what brings about the completion of the repeatable. Exceptions can occur, not through the laws failing but through additional factors not included in the available description of the situation.
3. *Factuality*. The continuable does in fact contribute to the completion of the repeatable. This is usually only assertible in retrospect, after the event in which the completable is completed, but can also occur in, for example, conditional contexts.

Once again, [2] has an operator *Prog* ready-made for this purpose. The operation in question may be called *progressivisation*. It maps a repeatable onto a continuable whose defining characteristic is that it is directed towards the completion of the continuable in one of the ways specified above. The continuables obtained in this way may be called *derived continuables*. The bottom row of Figure 5 shows the mappings related to the statement "Mary is singing *Happy Birthday*".

5 Processes revisited

In none of this discussion has the word "process" been used. How, then, do processes fit into the scheme expounded here? I propose that "process" should be regarded as a generic term that applies to all kinds of continuables and repeatables. One of the sources of confusion surrounding processes is that the term is applied to so many apparently different kinds of things. In particular, there is the use, widely assumed in linguistics, according to which, in the words of Mourelatos [10], "process" is "the topic-neutral counterpart of activity" (here referring to the well-known classification by Vendler [14] which forms the basis for the classification of perdurants in DOLCE). This corresponds well to our simple generic continuables, the things denoted by simple verbs such as "run", "sing", or "whistle" (the "topic-neutrality" being shown by an indifference as to whether it is a person or a river that runs, or a person or a kettle that whistles). These are, as already stated, homogeneous and open-ended. On the other hand, there is also the widespread use of "process" to refer to a completable routine comprising a structured sequence of actions or events. Thus we talk of the processes of

- ... making a pot of tea
- ... checking in for a flight
- ... assembling a model from a kit
- ... applying for a new passport.

These are the kinds of process that are typically referred to in expressions of the form "I am in the process of *X*ing". They correspond closely to our compound specific repeatables — compound because it is through composition that these processes can be specified as the result of particular ways of combining simple repeatables.

On the view taken here, then, processes are in the "abstract" realm. They are neither continuants nor occurrents, but rather can be realised as continuants — specifically states — or occurrents (events). This is very much in accord with the view expressed in [7, 11], and also fits in well with the abstract view of processes taken in computer science (see e.g., [8]). If this is right, then the relation between processes on the one hand and states and events on the other is quite different from any that are presented in the schemes illustrated in Figure 1. The category of process is neither subordinate to nor superordinate to the categories of state and event; nor is it on the same footing as them with some immediate common superordinate category. On the contrary, processes belong in a completely different realm, the realm of abstract entities, patterns if you will, quite separate from the realm of spatio-temporal entities which includes both states and events.

There remains a considerable amount of unfinished business. My passing reference to granularity in §3 needs to be amplified: what counts as a homogeneous process when

viewed at one level of granularity may appear to be structured at a finer granularity. A typical structured process consists of the periodic repetition of some sequence of events (e.g., the process of walking consist of a regular alternation of bounded leg-movements). This points to the need for more operators to allow the construction of these "higher-order" processes from events, and that in turn calls for a more detailed examination of the way in which events may be composed. A further omission here is any consideration of the category of instantaneous events (including Vendler's category of "achievements" [14]), and how these can be fitted into the proposed scheme. There is no space to go into these matters here; they will be handled in a future paper.

References

1. J. F. Allen. Towards a general theory of action and time. *Artificial Intelligence*, 23:123–54, 1984.
2. A. P. Galton. *The Logic of Aspect: an Axiomatic Approach*. Clarendon Press, Oxford, 1984.
3. A. P. Galton. Experience and history: Processes and their relation to events. *Journal of Logic and Computation*, 18(3):323–340, 2007.
4. A. P. Galton and R. Mizoguchi. The water falls but the waterfall does not fall: New perspectives on objects, processes and events. *Applied Ontology*, 4(2):71–107, 2009.
5. Pierre Grenon. Spatio-temporality in basic formal ontology: SNAP-SPAN, upper-level ontology, and framework for formalization, Part I. Technical Report 05/2003, Universitšat Leipzig Institute for Formal Ontology and Medical Information Science, 2003.
6. Pierre Grenon and Barry Smith. SNAP and SPAN: Towards dynamic spatial ontology. *Spatial Cognition and Computation*, 4(1):69–104, 2004.
7. Michael Grüninger. Using the PSL ontology. In S. Staab and R. Studer, editors, *Handbook on Ontologies*, International Handbooks on Information Syetems, pages 423–443. Springer, Heidelberg, 2009. Second edition.
8. C. A. R. Hoare. Communicating sequential processes. *Communications of the ACM*, 21(8):666–677, 1978.
9. M. Moens and M. Steedman. Temporal ontology and temporal reference. *Computational Linguistics*, 14:15–28, 1988.
10. A. P. D. Mourelatos. Events, processes, and states. In Philip Tedeschi and Annie Zaenen, editors, *Tense and Aspect*, pages 191–212. Academic Press, New York, 1981.
11. A. Özgövde and M. Grüninger. Functional process relations in bio-ontologies. In A. Galton and R. Mizoguchi, editors, *Formal Ontology in Information Systems: Proceedings of FOIS 2010*, pages 243–256, Amsterdam, 2010. IOS Press.
12. J. Pustejovsky. The syntax of event structure. *Cognition*, 41:47–81, 1991.
13. John F. Sowa. *Knowledge Representation: Logical, Philosophical, and Computational Foundations*. Brooks/Cole, Pacific Grove, CA, 2000.
14. Zeno Vendler. Verbs and times. In *Linguistics and Philosophy*, chapter 4. Cornell University Press, Ithaca, 1967. Reprinted in Mani, I., Pustejovsky J, and Gaizauskas, R, *The Language of Time: A Reader*, Oxford, 2005.