

Modelling Gradients Using Petri Nets

Laura M.F. Bertens¹, Jetty Kleijn², Maciej Koutny³, and Fons J. Verbeek¹

¹ Imaging and Bioinformatics group, LIACS, Leiden University
Leiden, 2300 RA The Netherlands
`lbertens@liacs.nl`, `fverbeek@liacs.nl`

² LIACS, Leiden University
Leiden, 2300 RA The Netherlands
`kleijn@liacs.nl`

³ School of Computing Science, Newcastle University
Newcastle upon Tyne, United Kingdom
`maciej.koutny@ncl.ac.uk`

Abstract. Motivated by the graded posteriorization during the AP axis development in the frog *Xenopus laevis*, we propose an abstract Petri net model for the formation of a gradient of proteins in a chain of cells.

Keywords: gradient formation, planar signalling, Petri net model

1 Introduction

Petri nets have been shown to be very promising for molecular and cellular biology, in particular for metabolic, signalling and gene-regulatory networks (see e.g. [1, 2, 6, 9, 10, 14, 20, 30, 31]). In this paper we propose Petri nets as an abstract modelling tool for higher level developmental processes in the organism, e.g., on tissue and organ level, taking cells as central elements.

Currently, we are working on a case study: the embryonic development of the anterior-posterior i.e., head-to-tail axis (AP axis) in the model organism *Xenopus laevis*, the African clawed frog. The development of this model embryo has been studied thoroughly and a huge amount of literature is available to draw from when building and validating the model, see references in [3, 15]. Moreover, this case study comprises several different subprocesses, found in many biological processes, that require modelling solutions. The aim of our project is to eventually model the entire process of AP axis development. Hence we envisage a final model consisting of several building blocks, most of which describing generic biological processes. Each of the subprocesses poses modelling challenges which, when solved, may lead to templates for similar developmental processes, incorporating multiple levels of both spatial and temporal information, also in other organisms. Petri nets are particularly useful in modelling such biological processes, due to their intuitive graphical component, which resembles biological diagrams, and their ability to model concurrency. Our case study appears to be very well suited to explore new ways in which Petri nets can be applied to developmental biology.

In this paper we present a fundamental approach to modelling a particular subprocess: the *formation of a morphogen gradient*, which helps instigate the differentiation of the cells along the developing axis. This subprocess is a good starting point, since it is relatively simple conceptually, in comparison to the other subprocesses in the case study. In early development, gradients are crucial ([36]) and finding a modelling solution for the generic process of gradient formation will not only serve the modelling of this case study, but will also be useful for the modelling of other developmental processes. By staying very close to the biological sequence of events in gradient formation, rather than focusing on a concrete outcome, the model should be generally applicable and robust.

Throughout this paper the emphasis will be on abstraction and modelling decisions, as opposed to implementation of specific biological data; we present a basic Petri net modelling gradient formation, which serves as a proof of concept for our approach. In the remainder of this paper we outline the biological background of gradient formation in general and in this particular case study. Subsequently we describe our modelling decisions and we present the model. In the last section the possibilities of the model and future work are discussed.

2 PT-nets with activator arcs

For a general introduction to Petri nets we refer to [27]. In this paper, we use PT-nets with activator arcs ([17]), and a maximally concurrent execution rule [5].

Petri nets are defined by an underlying structure consisting of *places* and *transitions*. These basic elements are connected by directed, *weighted arcs*. In the Petri net model considered in this paper, there are moreover *activator arcs* connecting places to transitions. In modelling, places are usually the passive elements, representing local states, and transitions the active elements. Here, global states, referred to as *markings*, are defined as mappings assigning to each place a natural number (of *tokens* corresponding to available resources).

A *PTA-net*, is a tuple $N = (P, T, W, Act, m_0)$ such that:

- P and T are finite disjoint sets, of the *places* and *transitions* of N , resp.
- $W : (T \times P) \cup (P \times T) \rightarrow \mathbb{N}$ is the *weight function* of N .
- $Act \subseteq P \times T$ is the set of *activator arcs* of N .
- $m_0 : P \rightarrow \mathbb{N}$ is the *initial marking* of N .

In diagrams, places are drawn as circles, and transitions as boxes. Activator arcs are indicated by black-dot arrowheads. If $W(x, y) \geq 1$, then (x, y) is an *arc* leading from x to y ; it is annotated with its weight if this is greater than one. A marking m is represented by drawing in each place p exactly $m(p)$ tokens as small black dots. We assume that each transition t has at least one input place (there is at least one place p such that $W(p, t) \geq 1$).

When a single transition t occurs (‘fires’) at a marking, it takes tokens from its input places and adds tokens to its output places (with the number of tokens consumed/produced given by the weights of the relevant arcs). Moreover, if there

is an activator arc $(p, t) \in Act$, then transition t can only be executed at the given marking if p contains at least one token, without the implication of tokens in p being consumed or produced when t occurs. Thus, the difference with a *self-loop*, i.e., an arc from p to t and vice versa, is that the activator arc only tests for the presence of tokens in p .

We define the executions of N in the more general terms of simultaneously occurring transitions. A *step* is a multiset of transitions $U : T \rightarrow \mathbb{N}$. Thus $U(t)$ specifies how many times transition t occurs in U . (Note that if we exclude the empty multiset, single transitions can be considered as minimal steps.) Step U is *enabled* (to occur) at a marking m if m assigns enough tokens to each place for all occurrences of transitions in U and, moreover, all places tested through an activator arc by a transition in U , contain at least one token.

Formally, step U is enabled at marking m of N if, for all $p \in P$:

- $m(p) \geq \sum_{t \in T} U(t) \cdot W(p, t)$
- $m(p) \geq 1$ whenever there is a transition t such that $U(t) \geq 1$ and $(p, t) \in Act$.

If U is enabled at m , it can be *executed* leading to the marking m' obtained from m through the accumulated effect of all transition occurrences in U :

- $m'(p) = \sum_{t \in T} U(t) \cdot (W(t, p) - W(p, t))$ for all $p \in P$.

Finally, a step U is said to be *max-enabled* at m if it is enabled at m and there is no step U' that strictly contains U (meaning that $U' \neq U$ and $U(t) \leq U'(t)$ for all transitions t) and which is also enabled at m . We denote this by $m[U]m'$. A (max-enabled) *step sequence* is then a sequence $\sigma = U_1 \dots U_n$ of non-empty steps U_i such that $m_0 [U_1] m_1 \dots m_{n-1} [U_n] m_n$, for some markings m_1, \dots, m_n of N . Then m_n is said to be a *reachable* marking of N (under the maximally concurrent step semantics).

To conclude this preliminary section, we elaborate a bit on the choice of this particular net model. First, it should be observed that it follows from the above definitions that the semantics allows *auto-concurrency*, the phenomenon that a transition may be executed concurrently with itself. This approach makes it possible to use transitions for a faithful modeling of natural events like the independent (non-sequential) occurrence in vast numbers of a biochemical reaction in a living cell. Note that the degree of auto-concurrency of a transition can easily be controlled by a dedicated place with a fixed, say k , number of tokens connected by a self-loop with that transition implying that never more than k copies of that transition can fire simultaneously.

Activator arcs were introduced in [16] as a means of *testing* for the presence of at least one token in a place, and so they are similar to other kinds of net features designed for the same reason. We mentioned already self-loops by which the presence of a token in a place can be tested only by a single transition (which ‘takes and returns’ the token) and not simultaneously by an arbitrary number of transition occurrences in a step. Two other mechanisms which do allow such multiple testing are *context* arcs [25] and *read (or test)* arcs [34]. Both, however, display important differences when compared with activator arcs. A context arc

testing for the presence of a token in place p by transition t indicates that after a step in which t participates has been executed, p must still contain a token which precludes the occurrence in the same step of transitions that have p as an output place. A read arc is also different, but less demanding in that there must exist a way to execute sequentially (i.e., one-by-one) all transition occurrences in the step, without violating the read arc specification. In both cases, one can easily see that activator arcs are most permissive since they only check for the presence of a token *before* the step is executed (this is often referred to as *a priori* testing). We feel that *a priori* testing is more appropriate for biological applications as the ‘lookahead’ implied by the other two kinds of test arcs is hard to imagine in reality.

Finally, we rely in this paper on *maximal* concurrency in the steps that are executed which reflects the idea that execution of transitions is never delayed. This may also be viewed as a version of time-dependent Petri nets where all transitions have a firing duration of 1. However, the maximal concurrency we apply in this paper does not derive from Petri nets with time, but rather from Petri nets with *localities* [19] leading to a *locally maximal* semantics. This semantics is what we plan to use to model other aspects of the development as well. Here one may think of e.g., the locally synchronous occurrence (in pulses) of reactions in individual compartments of a cell.

3 Biological background and modelling decisions

In biology, the term gradient is used to describe a gradual and directed change in concentration of a morphogen through a group of cells, e.g., a tissue. Morphogens are signalling molecules that cause cells in different places in the body to adopt different fates and thereby help establish embryonic axes. Morphogens are produced in a localized source of a tissue, the source cell(s), and emanate from this region, forming a concentration gradient ([13, 32]). A morphogen gradient has an immediate effect on the differentiation of the cells along it; cells are able to ‘read’ their position along the gradient and determine their developmental fate accordingly. They have a range of possible responses and the morphogen concentration dictates which response will be exhibited ([13, 32]). In establishing their developmental fate, cells take into account the morphogen concentration. When the morphogen concentration over the entire gradient is increased (or decreased), the cells should accordingly change their response to that corresponding to a higher (or lower) level of morphogen.

The mechanisms by which the morphogen travels through a cell layer have been the topic of some debate and are not yet fully understood. Three mechanisms have been described, shown schematically in Figure 1: (A) diffusion through the extracellular matrix ([8, 11, 23]), either passively, like a drop of ink in water ([11]), or facilitated by receptors on the cell surface which guide the morphogens along ([8]), as shown in the figure; (B) sequential internalization of the morphogen molecules in vesicles in the cells, a process called endocytosis, and subsequent re-emission ([7, 8, 32]); (C) direct contact between the cells by




Fig. 1. Left: three possible mechanisms for gradient formation: diffusion (A), endocytosis and subsequent re-emission (B) and transport through cytonemes (C). Right: an overview of the process of AP axis formation in *Xenopus laevis*

means of tentacle-like threads of cytoplasm, called cytonemes, connecting the cells ([13]). These mechanisms are not necessarily mutually exclusive and some studies conclude that a combination of mechanisms underlies the formation of a gradient. It is important to note that both diffusion and endocytosis take place between neighbouring cells, while cytonemes connect all cells directly to the source. This makes it very different from a modelling perspective, as will be discussed below.

Unfortunately, knowledge of the exact concentrations and shapes of most gradients is limited. This is mainly due to the transient nature of morphogen gradients and the low concentrations at which they are effective, both of which make it difficult to visualize the morphogens ([11]). Many morphogens are rapidly degraded or prevented from binding to receptors by antagonistic proteins ([11]). Much of the information on gradients is therefore obtained indirectly, by observing their effect, i.e., the responses of the cells along it ([11]).

AP axis formation in *Xenopus laevis*: a case study. The AP axis formation in *Xenopus laevis* takes place during the early embryonic stage of gastrulation and ensures the development of anterior structures near the head and posterior structures towards the tail. The process can be seen as divided into two steps, which take place sequentially ([26]), cf. Figure 1. The first step is activation; a group of cells in the outer cell layer of the embryo, the ectoderm, change their developmental identity and form a rectangular strip of tissue, called the neurectoderm ([3, 15]). It is this strip of cells in which the AP axis will ultimately be established, leading to gradual posteriorization of cells nearer to the tail-end of the embryo.

During the second step, transformation, the axis is formed in the neurectoderm by means of two mechanisms: vertical and planar signalling between neighbouring cells ([3, 15]). Here we focus on the second. Planar signalling occurs within the neurectoderm in a direction parallel to the future axis (and is therefore called ‘planar’). Concentration gradients of several morphogens are formed in the neurectoderm along the future AP axis. Source cells on the posterior end of the neurectoderm produce the morphogens, which then get distributed throughout the tissue. Individual cells sense their position along these gradients and take on a more or less posterior fate according to the concentration of these posteriorizing molecules ([12, 15]), thereby establishing the formation of an AP axis. In the planar signalling of our case study three types of signalling molecules play a role: retinoic acid (RA), fibroblast growth factors (Fgf) and Wnts ([22, 29, 35]). All of these are produced at the posterior end of the embryo and together these promote posterior cell fates, while inhibiting anterior fates. Although it is clear that all three types are important in axis formation, it is not yet fully understood how these proteins interact in establishing cell fate. A general and abstract modelling approach, focusing on the underlying common process of gradient formation, makes it possible to later add specific data on any of the morphogens in particular or on combinations of these.

Modelling decisions. We have chosen cells as the elementary units in our model to be represented as places in a Petri net. Earlier studies ([4, 21, 24]) have successfully modelled cell-to-cell signalling, starting from a lower biological level, using places to represent genes and proteins. Although this allows a high level of detail, it also complicates the net and makes it difficult to identify single cells. In our approach the cellular level represents the intermediate level between the subcellular levels, on which the morphogen signalling between cells takes place, and the tissue/organ level, where whole cell layers may move.

Tokens are used to represent a certain level of concentration (see [10]) within the overall gradient of the morphogen system, without differentiation between morphogens or their quantities. As mentioned before, in most cases no quantitative data are available, since morphogen gradients are often transient and difficult to visualize ([13]). As in [21], our approach is therefore partially qualitative and partially quantitative. The significance of tokens in a place is not purely qualitative; not only the presence or absence but also the exact number of tokens determine the course of events. However, the numbers of tokens do not represent actual numbers of molecules, making the model semi-qualitative. Our Petri net model can, however, also be used to model specific morphogens, incorporating quantitative data, by assigning exact concentrations to the tokens and thereby making the model completely quantitative. For certain morphogens quantitative data exist, for instance for the gradient of Fgf8 in zebrafish, which can be seen to spread extracellularly through the processes of diffusion, endocytosis and degradation ([28]). Also, for some gradients found in biological processes, experimental data have enabled to deduce mathematical expressions, describing the

quantitative morphogen concentrations ([33]). When modelling these gradients, these formulas can be incorporated in the parameters of our Petri net model.

Neighbourhood communication. It is our aim to develop a faithful model for gradient formation. Rather than having the net simply distribute the proper pre-computed amount of tokens over the places representing the cells, the actual transport of morphogen between cells can be read off from the Petri net model during execution. Consequently, when building the model we have to specify explicitly which process of gradient formation is to be modelled. Here we choose to model morphogens moving between neighbouring cells, i.e., the Petri net will implement a mechanism similar to diffusion or endocytosis and subsequent re-emission, but not transport through cytonemes (since this does not take place between neighbouring cells). However, we foresee no problems in the abstract implementation of the latter process. The difference between diffusion and endocytosis is apparent on a lower biological level and could be modeled by subnets. Furthermore, often the ratio of the concentrations between neighbouring cells is not known due to lack of quantitative data, and it may vary depending on the gradient considered. Therefore we have a parameter ρ in our model to represent this ratio and to determine the amount of tokens to be transported between places during the simulation of gradient formation. Since we do not distinguish the molecular mechanisms of diffusion, endocytosis and degradation of morphogens in this model, ρ represents the final ratio of morphogens between neighbouring cells and morphogen degradation is implicit. To model explicitly both the production of morphogens in the local source cells and the degradation in the target cells, subnets could be added. This should make the source and sink mechanisms of gradient formation transparent and allow the user to experiment with different configurations.

Implementation. In the organism, gradient ratios arise passively as a consequence of physical laws. However, to accurately reflect the biological process of gradient formation underlying the spread of morphogens from cell to cell, our formal model has to compute the number of tokens passed on based on the ratio ρ . Hence the model includes explicit separate computational units for the necessary calculations. In particular, these parts of the net control the transport of tokens between places. In this way a close relation to the biological process can be maintained in one part of the net, with the underlying computations performed by a subnet in the background. At all times, the marking of the net will be consistent with biological observations of (the effect of) the gradient, i.e., the ratio is maintained and places corresponding to cells further away from the source will never have more tokens than places (cells) closer to it. Another important feature of the model is the use of concurrent steps rather than individually occurring transitions. Cells only react to their environment and have no knowledge of other cells than their immediate neighbours. Non-adjacent cells can be simultaneously involved in the transport of morphogens. This leads to an execution mode consisting of concurrent *steps*. Moreover, these steps are *maxi-*

mal to reflect that also in the net model morphogens are moved to a next cell as soon as possible.

4 Gradients and Petri Nets

Following the ideas outlined in the previous section, we will propose a formal model for the formation of a gradient.

Our assumptions regarding the biological process of gradient formation are as follows. Given is a segment of k adjacent cells with the i -th cell immediate neighbour of the $(i + 1)$ -th cell. Morphogens can be transported only between immediate neighbours. Morphogens move from cells with higher concentration to neighbours with lower concentration, as long as their concentration ratio does not exceed a given gradient ratio $0 < \rho < 1$. We assume that ρ is a rational number, i.e., $\rho = \frac{N}{M}$, where $M > N \geq 1$. Initially, the first cell x_1 contains a quantity (has concentration level) K of a morphogen. These assumptions lead to the following modelling problem.

Given are $k \geq 1$ places x_1, \dots, x_k , representing a segment of k cells with place x_i corresponding to the i -th cell. In the initial marking m_0 , the first place x_1 contains K tokens and there are no tokens in the other places.

In the net modelling the mechanism of gradient formation, we need to shift tokens from x_1 in the direction of the last place x_k . Places and/or transitions may be added, but in such a way that for any reachable marking m the following hold.

1. The number of tokens in the x_i 's remains constant, i.e.,

$$m(x_1) + \dots + m(x_k) = K \quad \text{token preservation}$$

2. The tokens are distributed monotonically along the sequence of k places, i.e.,

$$m(x_1) \geq \dots \geq m(x_k) \quad \text{monotonicity}$$

3. The ratio of the numbers of tokens in two neighbouring places does not exceed ρ , i.e., for every $1 \leq i < k$ with $m(x_i) \geq 1$:

$$\frac{m(x_{i+1})}{m(x_i)} \leq \rho \quad \text{ratio}$$

4. Shifting continues until moving even one token would violate the above, i.e., if no tokens are shifted after marking m was reached, then for every $1 \leq i < k$ with $m(x_i) > 1$:

$$\frac{m(x_{i+1})+1}{m(x_i)-1} > \rho \quad \text{termination}$$

Moreover, the relative position of a place within the sequence plays no role. In particular, the mechanism should be easily scalable and **insensitive** to the specific values of k and K . \square

If we look at the above formulation of properties (2) and (3) — monotonicity and preservation of the gradient ratio — and recall that $\rho = \frac{N}{M}$ and $M > N$, it is easy to observe that these two properties are together equivalent to stating that, for every $1 \leq i < k$, $N \cdot m(x_i) - M \cdot m(x_{i+1}) \geq 0$. We will call a marking m satisfying this inequality *consistent* and denote $\alpha_i \stackrel{\text{df}}{=} N \cdot m(x_i) - M \cdot m(x_{i+1})$, for every $1 \leq i < k$. Note that the initial marking is consistent.

Similarly, if we look at the above formulation of properties (2) and (4) — monotonicity and termination — it is easy to observe that together they are equivalent to the statement that, for every $1 \leq i < k$, $N \cdot m(x_i) - M \cdot m(x_{i+1}) < M + N$. We will call a consistent marking m satisfying this inequality *stable*. Note that for a given ρ , k and K , there may be more than one stable marking. For example, if $\rho = \frac{1}{2}$, $k = 5$ and $K = 111$, then the following are two different stable markings:

$$\begin{array}{ccccc} x_1 & x_2 & x_3 & x_4 & x_5 & & x_1 & x_2 & x_3 & x_4 & x_5 \\ 59 & 29 & 14 & 6 & 3 & & 58 & 29 & 14 & 7 & 3 \end{array}$$

We are now ready to propose a generic solution for the above problem.

For a given consistent marking m and each $1 \leq i < k$, move β_i tokens from x_i to x_{i+1} where $\beta_i \leq \left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$, and at least one β_i must be non-zero if at least one of the values $\left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$ is non-zero. We denote the resulting marking by $m_{\beta_1 \dots \beta_{k-1}}$.

An intuitive reason for proposing such a mechanism for shifting tokens is that the number of tokens in x_i that are ‘balanced’ by tokens in x_{i+1} is $\frac{M}{N} \cdot m(x_{i+1})$, because each token in x_{i+1} is equivalent to $\frac{M}{N}$ tokens in x_i . Hence there are $m(x_i) - \frac{M}{N} \cdot m(x_{i+1})$ unbalanced tokens in x_i . The ‘portion’ of each unbalanced token that could be safely transferred to x_{i+1} is $\frac{N}{M+N}$. Hence in total we may safely transfer $\left\lfloor \frac{N}{M+N} \cdot (m(x_i) - \frac{M}{N} \cdot m(x_{i+1})) \right\rfloor$ tokens, which is precisely $\left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$ tokens. Clearly, some of the numbers $\beta_1, \dots, \beta_{k-1}$ can be zero, and by the condition above, all β_i ’s are zeros if and only if the marking is stable:

Proposition 1. $\beta_1 = \dots = \beta_{k-1} = 0$ if and only if m is stable.

Crucially, by the mechanism proposed consistent markings are always transformed into consistent markings.

Proposition 2. If m is a consistent marking then $m_{\beta_1 \dots \beta_{k-1}}$ is also consistent.

According to the above, any number of tokens not exceeding $\left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$ can be moved *simultaneously* from x_i to x_{i+1} (for every $i < k$), and consistency will be preserved. Clearly, the new consistent marking is different from the previous one if and only if, for at least one i , we have $\beta_i \geq 1$. The idea now is to keep changing the marking on x_1, \dots, x_k until a marking m has been reached such that $\left\lfloor \frac{\alpha_i}{M+N} \right\rfloor = 0$, for all $1 \leq i < k$, which is equivalent to $\alpha_i < M + N$, for all $1 \leq i < k$. In other words, this m is a stable marking. Since tokens cannot

be shifted forever, this procedure will always terminate in a stable marking (formally, we can show this by considering a weighted distance to the end of the chain of the K tokens; it never increases and always decreases in a non-stable state).

Looking now from the point of view of a Petri net implementation of the proposed mechanism, what we are after is a net N_{shift} comprising the places x_1, \dots, x_k and such that if m is a marking of N_{shift} whose projection on these k places is consistent, then a step U can occur at m if

- it moves at most $\left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$ tokens from x_i to x_{i+1} , for all $1 \leq i < k$;
- at least one token is moved from x_i to x_{i+1} for at least one $1 \leq i < k$, unless the projection of m onto x_1, \dots, x_k is stable.

In fact, in the proposed implementation, we will be preceding the ‘token-shifting’ with a ‘pre-processing’ stage which seems to be unavoidable unless one uses some kind of arcs with complex weights depending on the current net marking.

Implementation. In the implementation of the proposed shifting mechanism, as many tokens as possible should be shifted from one neighbour to the next. That means that, at each stage we have, for every $1 \leq i < k$, $\beta_i = \left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$. Moreover, tokens are shifted from a place without any assumptions whether new tokens will come to that place from its other neighbour. Thus we need to provide a Petri net structure capable of ‘calculating’ the value of expressions like

$$\left\lfloor \frac{N \cdot m(x_i) - M \cdot m(x_{i+1})}{M+N} \right\rfloor.$$

Our proposed gradient forming mechanism distinguishes three phases: I, II and III. An auxiliary net N_{3phase} , shown in Figure 2(b), is used to schedule the transitions implementing the calculations. It controls these transitions via the places w^I and w^{II} and activator arcs. For the full picture of the system one should combine the figures for all pairs (x_i, x_{i+1}) with a single copy of the net in Figure 2(b). Note that all places with identical label (in particular w^I , w^{II} , and w^{III}) should be identified. That other parts of the encompassing net model do not interfere with the calculations carried out during phases I and II can be ensured by connecting the relevant transitions with the place w^{III} using activator arcs.

For every $1 \leq i < k$, transition t_i is intended to shift tokens from x_i to x_{i+1} (phase III). To achieve this, we use two disjoint sets of new, auxiliary places, x'_1, \dots, x'_k and x''_1, \dots, x''_k . These places are initially empty. The idea is to fill x'_i with $N \cdot m(x_i)$ tokens and x''_{i+1} with $M \cdot m(x_{i+1})$ tokens (phase I). The latter are used for the removal of $M \cdot m(x_{i+1})$ tokens from x'_i (phase II). After this, there are α_i tokens remaining in x'_i . Finally, for each group of $N + M$ tokens in x'_i , one token is shifted from x_i to x_{i+1} . The construction (for x_i and x_{i+1}) is shown in Figure 2(a).

The overall mechanism operates in cycles of three consecutive, maximally concurrent steps such that for every $1 \leq i < k$:




Fig. 2. (a) The main part of the construction for the solution (note that e''_{i+1} is introduced for later use when one might want to remove or add tokens to the x_i 's from 'outside'; in the standard (consistent) situation it is never activated as after phase 2, place x''_{i+1} is empty.); and (b) the subnet N_{3phase} enforcing the three phases.

- I. Transition c'_i , inserts (in $m(x_i)$ auto-concurrent occurrences) $N \cdot m(x_i)$ tokens into x'_i . In the same step, transition c''_{i+1} , inserts (in $m(x_{i+1})$ auto-concurrent occurrences) $M \cdot m(x_{i+1})$ tokens into x''_{i+1} . Simultaneously, transitions e'_i and e''_{i+1} empty x'_i and x''_{i+1} of any residual tokens left from the previous cycle.
- II. Next, transition d_i (in $M \cdot m(x_{i+1})$ auto-concurrent occurrences) empties x''_{i+1} and leaves in x'_i the difference $\alpha_i = N \cdot m(x_i) - M \cdot m(x_{i+1})$.
- III. In the third step, the occurrences of transition t_i transfer $\beta_i = \left\lfloor \frac{\alpha_i}{M+N} \right\rfloor$ tokens from x_i to x_{i+1} .

Proposition 3. Each cycle results in transferring β_i tokens from x_i to x_{i+1} .

Note that in this implementation with the control net N_{3phase} , neighbouring pairs are either all involved in calculations (step 1 and 2 of the cycle) or tokens are transferred between neighbours (step 3). During the whole operation of the adjustment process (except for the transfer phase), the token counts in the places x_i representing the cells are unchanged and they can be accessed for reading by other transitions (and thus influence neighbouring cells). In other words, calculations are *orthogonal* to the basic operation of the net (the gradient formation).

As an example, let us consider the case when $\rho = \frac{1}{2}$, $k = 4$ and $K = 100$. Then executing the constructed net in a maximally concurrent manner leads to the following sequence of markings on the x_i after each cycle and eventually to a stable marking:

x_1	100	67	67	60	60	57	57	56	56	55	54
x_2	0	33	22	29	25	28	26	27	26	25	26
x_3	0	0	11	8	12	10	12	12	13	12	12
x_4	0	0	0	3	3	5	5	5	5	6	6

The next example shows what happens if we start from a (non-initial) consistent marking (again $\rho = \frac{1}{2}$):

x_1	200	167	156	...
x_2	50	67	67	...
x_3	0	16	22	...
x_4	0	0	5	...

The construction works without any problems, if we start with a consistent marking. In case $0 > \alpha_i$ for some i , then transition t_i is not executed, but the transitions t_{i-1} and t_{i+1} may still be executed and lead to an adjustment of the marking causing t_i to become active in the next cycle. A further observation is that adding (or removing) tokens at some point, will trigger a re-adjustment process which tries to re-establish the correct ratios between the markings of adjacent places x_i . This process is unpredictable, but to deal with that case we have included transition e''_{i+1} which in the standard (consistent) situation is never activated since then, after phase 2, place x''_{i+1} is empty.

An important characteristics of the proposed solution is that it is purely *local* and does not assume anything about the number of tokens which may appear in the x_i 's nor the length of the chain. In other words, it is truly generic. What's more it also works if M and N are different for different pairs of neighbouring places, i.e., if rather than a uniform gradient ratio ρ there is a ratio ρ_i for each pair of neighbours x_i and x_{i+1} .

Another feature of our solution is the maximal concurrency semantics intended to reflect the idea of morphogens (simultaneously) moving from cell to neighbouring cell whenever that is possible. The preliminary sequential semantics model we developed (but not reproduced here) is more complicated as it also needs *inhibitor arcs* which test for *absence* of tokens (to decide whether or not tokens should still be shifted). Moreover, one needs to decide that x_i either receives or sends tokens at each stage. In a step model it can both receive and send. Also, with the maximal concurrency semantics, the number of states of the model is dramatically reduced.

The auxiliary net N_{3phase} is used to partly sequentialize the behaviour in order to separate the pre-processing phases from the actual shifting phase. This net could also have been made local to the subnet in Figure 2(a), with different copies of it assigned to different localities. This would have given the additional possibility of controlling the degree of synchronisation between different parts of the gradient model by using a locally maximal step semantics.

Finally, we would like to point out that the activator arcs in our implementation are used only to control the calculation and can actually be avoided in case there would be a limit on the number of tokens in each place x_i at any time. (Then the activator arcs can be eliminated basically by having separate copies of N_{3phase} for each $1 \leq i < k$, transfer around sufficiently many tokens in a bundle, and replace activator arcs by self-loops). This assumption corresponds to having (or knowing) some capacity bound on the concentration levels of morphogens in a cell and so may be biologically sound!

5 Conclusion

Starting from gradient formation in the AP axis development in the model organism *Xenopus laevis*, we have presented a novel approach to using Petri nets in developmental biology by focusing on the cellular rather than subcellular levels and abstracting from concrete proteins and genes. This has led to a parameterized Petri net model for the general process of gradient formation through diffusion and endocytosis.

Assumptions regarding gradient formation have been formulated based on essential features of this process as reported in the literature. These assumptions underlie the precise requirements given that should be satisfied by an abstract Petri net model of gradient formation. A crucial point here is the consistency that is maintained during the execution of the model. Hence the realization of the gradient is faithfully reflected. Moreover the close relationship between biological process and evolution of the formal model makes it possible to apply existing Petri net techniques to analyse what happens during gradient formation. In particular cause-effect relations should be properly reflected in the process semantics of the modelling Petri net ([17, 18]).

Another main contribution of this approach is its generic nature, leading to a model that is scalable and applicable to a plethora of specific gradients. Also scalability is a consequence of the faithful reflection of the biological process. Since the final token (morphogen) distribution is not directly computed from the initial amount of morphogen and the length of the chain of cells, but rather simulates the communication between neighbouring cells, the length plays no role in the occurrence of the steps. The model as presented here represents a one-morphogen system without relying on quantitative data, but exact values could be assigned to ratio and individual tokens. Moreover, it provides a basis for simulation of simultaneous gradient formation (different morphogens with different experimental initial markings) and for inhibiting/activating interactions between them. Simulation with actual biological data to validate the model should be a next step. In addition, we will focus our attention on the extension of this still rather basic model to more dimensions, e.g., rather than having just a single line of cells, we consider the spread of morphogens from a source throughout a tissue plane or volume.

In [4, 21], Petri nets are used to model developmental processes in a way similar to our approach when it comes to the semi-qualitative use of tokens and the use of maximal concurrency. In these papers however, the focus is on subcellular levels. Petri net places are used to represent genes and gene products, where in our approach cells, as basic units in a tissue, are modelled by places. Having cells as basic units should prove to be a useful intermediate position convenient for ‘zooming in and out’ between subcellular and tissue level. It is our aim to model more subprocesses of the AP axis formation. For instance the different molecular processes underlying diffusion and endocytosis could be modeled in subnets, allowing the user to compare the different effects of these mechanisms. Also the degradation of morphogens could be modeled by a subnet, making the entire process more explicit. The choice of cells as main elements is

expected to be particularly suitable not only in this ‘vertical’ linking processes, but also for the ‘horizontal’ connections between processes taking place on the same cellular level. Having molecules or genes as places would result in specific net models for certain processes, as would taking tissue structures, such as the neurectoderm. Cells however can play a role in different processes simultaneously.

The next step in the modelling of the AP axis development in *Xenopus laevis* will focus on the vertical signalling (see Figure 1). This process occurs concurrently with the planar signalling of gradient formation and involves the same cells. This will challenge us to explore further the possibilities of Petri nets as a model for concurrent and independent processes in high level developmental biology.

Acknowledgments We are grateful to all referees for their remarks and suggestions which helped us to improve the paper. We would like to express also our appreciation to A.Durstun and H.Jansen for discussing *Xenopus* biology.

References

1. R.Banks: Qualitatively Modelling Genetic Regulatory Networks: Petri Net Techniques and Tools. Ph.D. Thesis, Newcastle University, 2009
2. R.Banks, V.Khomenko, L.J.Steggles: A Case for Using Signal Transition Graphs for Analysing and Refining Genetic Networks. ENTCS 227, 2009, 3–19
3. N.Bardine: Hox in Frogs. PhD Thesis, University of Leiden, 2008
4. N.Bonzanni et. al: Executing Multicellular Differentiation: Quantitative Predictive Modelling of *C.elegans* Vulval Development. Bioinformatics 25, 2009
5. H.-D.Burkhard: On Priorities of Parallelism: Petri Nets under the Maximum Firing Strategy. LNCS 146, 1983, 86–97
6. C.Chaouiya: Petri Net Modelling of Biological Networks. Briefings in Bioinformatics 8, 2007, 210–219.
7. E.Entchev, M.Gonzalez-Gaitan: Morphogen Gradient Formation and Vesicular Trafficking. Traffic 3, 2002, 98–109
8. J.A.Fischer, S.H.Eun, B.T.Doolan: Endocytosis, Endosome Trafficking, and the Regulation of *Drosophila* Development. The Annual Review of Cell and Developmental Biology 22, 2006, 181–206
9. D.Gilbert, M.Heiner: From Petri Nets to Differential Equations - an Integrative Approach for Biochemical Network Analysis. LNCS 4024, 2006, 181–200
10. D.Gilbert, M.Heiner, S.Lehrack: A Unifying Framework for Modelling and Analysing Biochemical Pathways Using Petri Nets. LNBI 4695, 2007, 200–216
11. J.B.Gurdon et. al: Activin Signalling and Response to a Morphogen Gradient. Nature 371, 1994, 487–492
12. J.B.Gurdon et. al: Single Cells Can Sense their Position in a Morphogen Gradient. Development 126, 1999, 5309–5317
13. J.B.Gurdon, P.Bourillot: Morphogen Gradient Interpretation. Nature 413, 2001, 797–803
14. M.Heiner, D.Gilbert, R.Donaldson: Petri Nets for Systems and Synthetic Biology. LNCS 5016, 2008, 215–264
15. H.J.Jansen: Anterior-posterior Axis Formation in *Xenopus laevis*. PhD Thesis, University of Leiden, 2009

16. R.Janicki, M.Koutny: Semantics of Inhibitor Nets. *Information and Computation* 12, 1995, 1–16
17. J.Kleijn and M.Koutny: Processes of Petri Nets with Range Testing. *Fundamenta Informaticae* 80, 2007, 199–219
18. J.Kleijn and M.Koutny: Processes of Membrane Systems with Promoters and Inhibitors. *Theoretical Computer Science* 404, 2008, 112–126
19. J.Kleijn, M.Koutny, G.Rozenberg: Towards a Petri Net Semantics for Membrane Systems. *LNCS* 3850, 2006, 292–309
20. I.Koch, B.H.Junker, M.Heiner: Application of Petri Net Theory for Modelling and Validation of the Sucrose Breakdown Pathway in the Potato Tuber. *Bioinformatics* 21, 2004, 1219–1226
21. E.Krepska et. al: Design Issues for Qualitative Modelling of Biological Cells with Petri Nets. *LNBI* 5054, 2008, 48–62
22. T.Kudoh, S.W.Wilson, I.B.Dawid: Distinct roles for Fgf, Wnt and Retinoic Acid in Posteriorizing the Neural Ectoderm. *Development* 129, 2002, 4335–4346
23. A.D.Lander, Q.Nie, F.Y.M.Wan: Do Morphogen Gradients Arise by Diffusion? *Developmental cell* 2, 2002, 785–796
24. H.Matsuno et. al: Boundary Formation by Notch Signaling in Drosophila Multicellular Systems: Experimental Observations and Gene Network Modeling by Genomic Object Net. *Pacific Symposium on Biocomputing* 8, 2003, 152–163
25. U.Montanari, F.Rossi: Contextual Nets. *Acta Informatica* 32, 1995, 545–596
26. P.D.Nieuwkoop: Activation and Organisation of the Central Nervous System in Amphibians. III. Synthesis of a New Working Hypothesis. *Journal of Experimental Zoology* 120, 1952, 83–108
27. W.Reisig, G.Rozenberg (Eds.): Lectures on Petri Nets I and II. *LNCS* 1491 and 1492, 1998
28. S.Scholpp, M.Brand: Endocytosis Controls Spreading and Effective Signaling Range of Fgf8 Protein. *Current Biology* 14, 2004, 1834–1841
29. J.Shiotsuguet et. al: Multiple Points of Interaction between Retinoic Acid and Fgf Signaling During Embryonic Axis Formation. *Development* 131, 2004, 2653–2667
30. L.J.Steggles, R.Banks, O.Shaw et al.: Qualitatively Modelling and Analysing Genetic Regulatory Networks: a Petri Net Approach. *Bioinformatics* 23, 2006, 363–343
31. C.Talcott, D.L.Dill: Multiple Representations of Biological Processes. *Transactions on Computational Systems Biology*, 2006
32. A.A.Teleman, M.Strigini, S.M.Cohen: Shaping Morphogen Gradients. *Cell* 105, 2001, 559–562
33. C.J.Tomlin, J.D.Axelrod: Biology by Numbers: Mathematical Modelling in Developmental Biology. *Nature Reviews Genetics* 8, 2007, 331–340
34. W.Vogler: Partial Order Semantics and Read Arcs. *Theoretical Computer Science* 286, 2002, 33–63
35. R.J.White et. al: Complex Regulation of *cyp26a1* Creates a Robust Retinoic Acid Gradient in the Zebrafish Embryo. *Plos biology* 5, 2007, 2522–2533
36. L.Wolpert: *Principles of Development*. Oxford University Press, 2002