PO Box 1127
Rotorua 3040
Ph: + 64 7 921 1883
Fax: + 64 7 921 1020
Email: info@ffr.co.nz
Web: www.ffr.co.nz

Theme: Radiata Management

Task No: F10203 **Report No. : R069**
Milestone Number: 2.03.9

Biomechanical Framework Integration Test

Authors:
D Sellier, J Harrington

Research Provider:
Scion

Date: January 2012

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Forest managers and forest researchers share a need for better understanding of how trees function. At any point in time, the way trees operate determines both the volume and the quality of the wood resource being produced. There are two aspects to the problem: what are the mechanisms underlying the wood factory (trees), and what are the inputs and operational conditions (the environment). Dependency of wood formation on environmental variables is a well-known fact that makes trees some of the most reliable climate markers. In the scientific community, many mechanisms have been proposed to explain how the rate at which wood is being laid and how wood fibre characteristics are regulated by factors such as temperature, available water, atmospheric CO_2, and so on. So far, those mechanisms have been investigated independently of each other.

A modelling framework has been developed over years as part of Intensive Forest Systems objective 2 *'Mechanobiology of wood formation'*. The purpose of this framework is to federate and integrate the knowledge about tree physiology and growth biophysics. Many component models have been developed in order to simulate in a deterministic manner key aspects of tree function such as the transport of photoassimilates on tree stem surface.

In this project, we performed a simulation run which integrated all available framework components except those related to mechanics, which will be tested as part of a future milestone. The primary goal of this exercise was to identify challenges and issues for simulating tree growth at large scale. It is the first time a framework has combined all-at-once a process-based growth model, an architectural model, finite element analyses to simulate carbon transport, a process-based cambium model and also models for predicting wood properties.

The integration test showed it is possible to simulate tree growth and development as a complex system with interdependent components in a deterministic manner. Two important challenges lie ahead for the "Mechanobiology of wood formation" framework. The first challenge will be to have access to supercomputer resources to perform large scale or repetitive simulations. The second and probably biggest challenge will be to manage the increasing demand for software engineering while also using the framework for what it has been first designed for: exploring new scientific avenues.

INTRODUCTION

The tree is the wood factory. Trees in reality are much more than that, but they are mainly grown for their wood. Like human-made factories, a tree requires many resources and involves many processes by which resources are transformed into the final product. Unlike human-made enterprises, how the factory works is not readily understood. Furthermore, while the wood produced probably tends towards what's best for the tree at any point in time, it does not always meet expectations as a resource for technological use.

To a large extent, each process linked to wood formation can be described reasonably well in isolation. They have been the subject of scientific studies for decades. Photosynthesis is a good example: one can predict with accuracy the amount of carbon assimilated by a plant from the light environment and atmospheric CO_2 levels. One can use more elaborate models and include more factors, but the base process and the key factors are well documented. Of course, in reality, no process exists in isolation. Growth processes all impact each other, they are all interdependent, which makes things complex... but also interesting. To come back to the example of photosynthesis, CO_2 enters plant leaves through small pores. The same pores are used for respiration and transpiration – like the human body, plants absorb oxygen to produce metabolic energy and release water to avoid over-heating. The fact that those pores (called stomata) serve several purposes implies there is a trade-off as to when they can be opened or when they should be closed; for example, stomata may stay closed on a sunny day in near drought conditions to avoid losing water in spite of the abundance of light. This is just an example. Such trade-offs and links between growth mechanisms are numerous within a tree. Although necessary to appreciate wood variability and its causes, process interaction is an aspect that is rarely taken into account.

One of the reason process interaction is neglected is that monitoring one process, let alone many, in a growing organism without impacting on its development is not a trivial task. Another difficulty is that processes are coupled. When studying and altering one process, a chain of consequences is triggered throughout the organism. It is hard to identify cause/effect relationships in those conditions, especially when so-said relationships can be indirect, conditional, non linear or buffered. Numerical analysis can help. Simulation tools allow multiple mechanisms to be described and interactions to be specified. They can describe the behaviour in an integrated way, a behaviour which can be compared to real life behaviours. They can also be used to infer the relative contribution of factors on a key output, such as growth rate. For instance, Daudet *et al*.[1] published a coupled model of carbon and water flows in a plant. The model is based on generally accepted theories on how those flows physically operate. Both carbon and water flow in the plant through different pathways, but they occur simultaneously and the pathways are connected. To our best knowledge, this is the first time both flows have been modelled concurrently. Their study showed that the water flow had little influence on the carbon allocation. Therefore the study has been instrumental in showing that the coupling is loose. Carbon allocation models neglecting the water flow can continue to do so. It shows that integrating multiple processes and the associated complexity can lead to simplifications, and simplifications that have been demonstrated. If the study had shown that water flow played a major role, all models previously neglecting the water flow "because it's just simpler that way" would have been invalidated. Numerical models always allow refining or redefining the description of an individual process, for example to test a new hypothesis, and observe the impact on the system as a whole. As such, they are also a great tool to identify gaps and misconceptions.

In this study we tested the integration of several models simulating specific aspects of tree development: photosynthesis/respiration, photoassimilate transport, wood growth rate, wood density, fibre orientation… The purpose of the test was to evaluate how the framework behaves and how components interact on a technical basis. The validation of simulations against real tree growth data was out of the scope of this exercise. Also, because the modelling framework focuses on simulating wood formation, aspects not directly related to wood formation can be modelled in a less detailed manner.

At any point in time, the tissue that produces wood (the cambium, henceforth) creates cells whose characteristics depend on many influences. Well-known influences are water availability (Abe *et al*. 2003), mechanical stress[2], the photoperiod[3], auxin and carbohydrates[4], and temperature[5]. The combination and the variation of environmental inputs result in wood being a very heterogeneous material[6]. Genetics also contribute significantly to wood variability[7]. However, it can be argued that genotypes largely contribute to variability in that they condition responsiveness of growth processes to the environment variables.

The cambium is located under trees' bark. As new wood is laid out, the cambial surface expands. The modelling framework we used in this study emphasizes the view of a tree as "growing living skin". Most framework components are there to determine the local environment or surface state at any point on that growing skin. Once the surface state is known, a cambial model is used to predict the radial growth rate corresponding to the growth environment.

METHODS

Framework Structure

The modelling framework developed for IFS objective 2 “Mechanobiology of wood formation” is composed of multiple computer models and tools. These can be divided into three main categories:

- surface evolution,
- surface state prediction, and
- cambial activity.

Surface evolution is in fact only one component, but a central one. It derives from the view that a tree is a living skin, always growing. The surface evolver (SE) component takes the tree stem and branch surfaces at any given time, and produces a new surface that corresponds to the next time step. To perform that operation, a field of growth rates is required. Growth rates can vary all over the tree surface, with a spatial resolution down to a few millimetres. This allows highly localised variations of growth rate and also wood properties to be modelled. Performing such operations at the scale of the whole tree has obviously a high computational cost, and this point will be monitored during this study.

The second category contains components that are used to define the state of the cambial surface. State is expressed as a set of environmental variables. Those variables include sucrose concentration, water potential, mechanical stress and temperature. Any variable can be defined provided that a component model can compute its value at any point of the surface, at any time during growth. Most components in that category simulate the distribution of a variable over the surface. A few others have an indirect role and exist only to generate input for the distribution models. For instance, we use an ecophysiological model in the framework to predict the amount of photoassimilates in the crown. The information is then passed to the component in charge to simulate the transport of assimilates (carbon) in the tree.

The third type of components relates to cambial activity. There is predominantly one component that returns a rate of radial expansion associated to the local growth environment. That component is the cambial model. It encapsulates the growth behaviour of the cambium. In the framework, several cambial models can be defined to test different hypotheses on the wood formation mechanisms. However, only one cambial model can be used for the duration of one simulation. The surface evolver is independent of which cambial model is used, in that it needs the growth velocity field output only. Other components linked to cambial activity can be included in the framework. They also take the local environment as input, but predict wood properties rather than the growth rate. Again, the framework is flexible. Many models of variable complexity can be incorporated. Only wood density and fibre orientation are modelled in this exercise.

Simulating the Growing Tree – Overview

Figure 1 shows the main steps in simulating the growing tree with the biomechanical framework. A tree of age 0 is created at initialisation. At this stage, some data are pre-computed and will be available for the rest of the simulation. The simulation then enters the main loop, which is repeated for each growth cycle until the simulation’s end time is met. The end time can be defined by the user. By default, it is the last age of the tree for which growth data are available. At each cycle, primary growth is simulated by inserting new surface elements into the cambial surface. Those elements are either added to existing shoots to reflect the elongation process, or inserted laterally when branching. The second step of the main loop determines the state of the cambial surface.. In other terms, this is the stage where all variables defining the local growth environment are computed. The third step consists of calculating the local growth rate that matches the local environment. It is done at all points of the surface with a resolution of a few millimetres. It is thus possible to describe not only longitudinal variations in the growth rate but also circumferential variations, as arise, for example, when compression wood is formed. The fourth and final step in the main loop is to evolve the surface geometry as a function of the growth rate field determined during the previous step. Wood properties are also evaluated during this last stage. Then the simulation routine returns to the beginning of the loop and reiterates the four-step sequence for the next growth cycle.

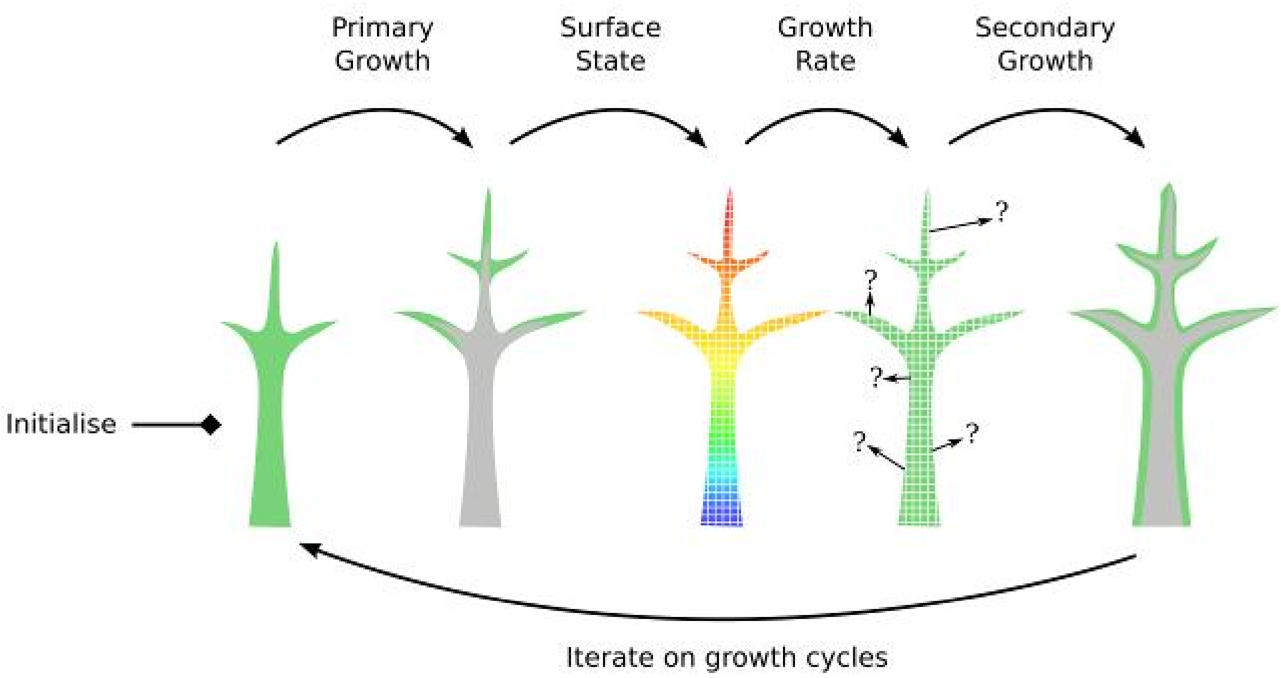

Figure 1: Main steps of tree growth simulation

Modification of Individual Components

We briefly describe how the framework components were modified in order to carry out the integration test.

Surface Evolver

The surface evolver is the cornerstone of the biomechanical framework. As a consequence, it is the oldest and most tested component. However, further development was necessary to carry out the present test. The level set method[8], a mathematical tool used here to simulate surface growth, implies using a computational grid one dimension higher than the interface being modelled. In our case, we perform calculations on a small volume wrapping the tree stem and branch surfaces. The problem is that the growth rate is predicted only on the surface by the cambial model, whereas the SE needs it defined everywhere in the volume. The solution is to employ an extrapolation method that propagates growth rate information from the surface throughout the volume. The extrapolation method we have chosen to implement is called extension velocities[9]. Adaptation of the seminal work was necessary to handle a multi-compartment surface (in the framework, the stem and each branch may follow their own growth patterns).

Cambial Model

The model used to simulate cambial growth is that of Deleuze and Houllier[10]. The model has been implemented as a framework component as part of previous milestone (F10203 – 11/12 – ms1). Minor modifications were done to integrate the model. Those modifications consist of unit conversions and replacing global quantities (stand carbon production, soil moisture) by local ones (sucrose concentration, water potential). The global to local change makes the Deleuze and Houllier's model similar to that of Holtta *et al.*[11].

Wood Properties

Two components predict macroscopic wood properties in this integration test: one component for basic density and one for grain orientation. Others could be included but those two are essential. Firstly, density is a scalar field and grain is a vector field. Thus, both encompass most mathematical types that one would want to use to store wood properties. Secondly, density is crucial for the relationships between mass and volume, while grain direction allows the material orthotropy of the wood material to be taken into account. For instance, the grain vector is used by the phloem transport model to simulate carbon allocation. The underlying models for density and grain are simplistic but sufficient for what is essentially a data handling exercise.

Phloem Transport

This component has been described in a previous technical note and was not heavily modified in this study. It is noteworthy that the component depends on finite element software. As the cambial surface grows, computations becomes intensive time-wise and memory-wise. Those computations are executed remotely on a more capable computer. Results of transport analysis are automatically transferred and made available to other components of the framework.

Xylem Transport

The xylem transport model works and operates similarly to the phloem transport model. It uses finite element analysis and pressure-based flow equations to simulate sap movement. The component exists and has already been presented to FFR. However, it has not been possible to include this component the integration test as it does not match the level of development of the other components. Mostly, the component does not handle material orthotropy or, in simpler words, cannot include grain orientation. It would also slow down considerably the integration test because it has higher computational requirements than the phloem transport model as it operates on the tree volume, not just the surface.

Growth Biomechanics

The role of mechanical state in cambium development will be investigated in the next milestone. Growth biomechanics were not considered in this study.

Architecture

The phloem transport model requires information about the weight of foliage borne by branches. In order to predict it, a relationship between basal area increment and foliage weight has been added in the cabalaTree component. Foliage weight is then stored as attribute data in the graph encoding the architecture. Another modification done to the architectural model was to add the ability to re-sample the architectural graph so that primary growth can be simulated with arbitrary time steps instead of yearly only, as originally prescribed.

Cabala

The Cabala component is a process-based growth model developed at CSIRO (Battaglia *et al*. 2004). The model itself was not modified for this study. However the temperature information is now extracted from the growth scenario results. The temperature is used by the cambial model to determine the number of cell divisions taking place at each time step. The temperature is considered homogeneous on the cambial surface. While vertical temperature gradients can be expected in the air surrounding the tree, sap convection within tree stems is expected to homogenise temperature distribution.

Component Interactions

Figure 2 shows how the framework components were connected for carrying out the integration test. It also shows the four different data types used to represent the tree:

- The MTG description, used for representing tree architecture;
- Time series of growth data for global or tree-wide variables;
- The surface, a triangulated surface describing the vascular cambium of the tree at a given time. Attribute data can be attached to any point or triangle of the surface
- The wood formed, which is the volume contained within the surface. It is where characteristics of wood fibres at the time of formation are stored.

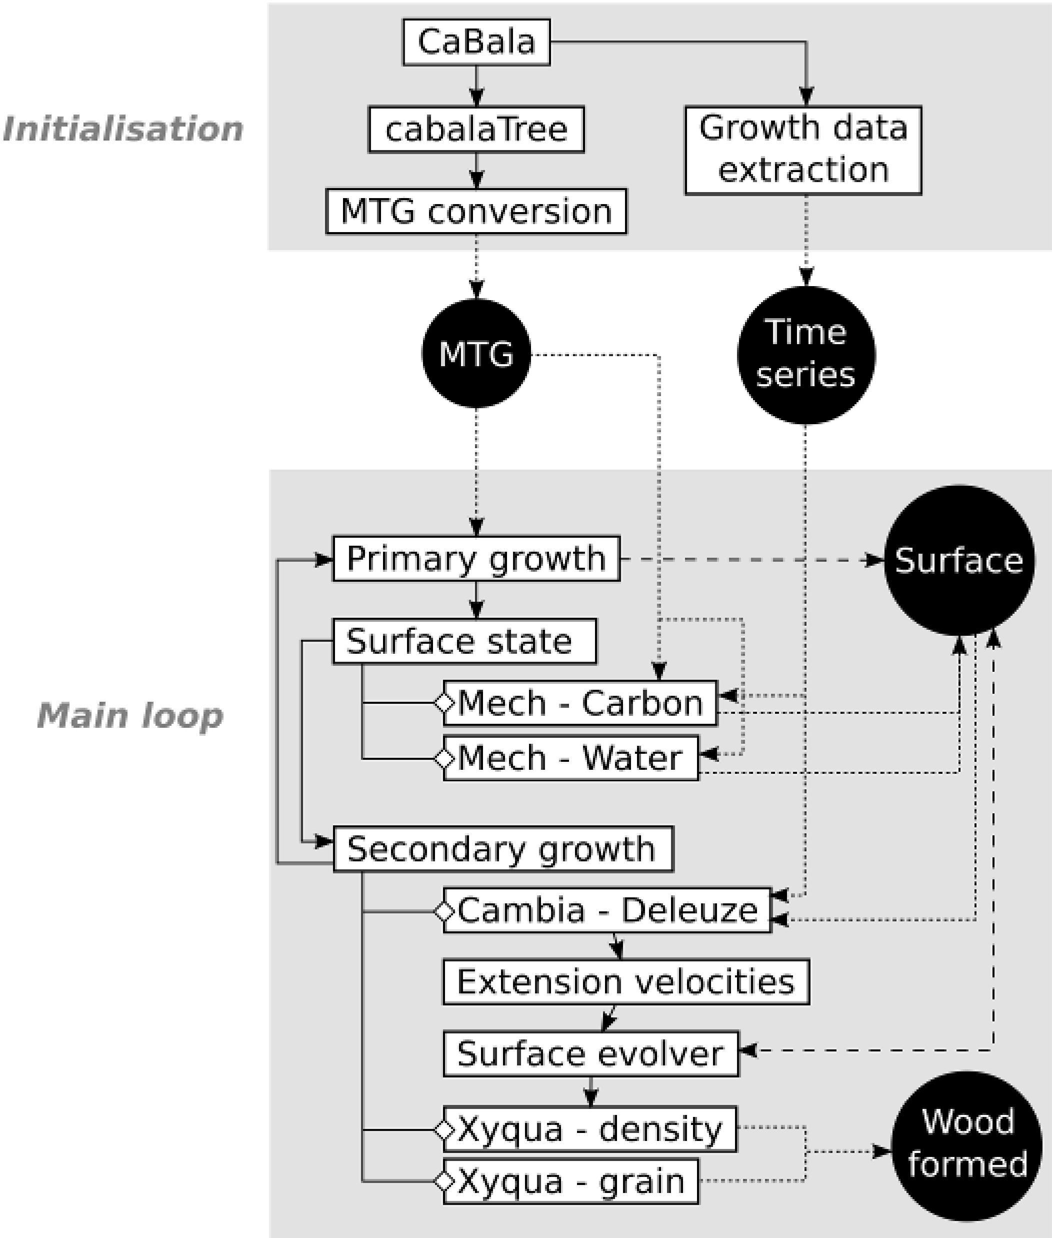

Figure 2: Component interaction and tree data representation.

Components (boxed names) in the initialisation group are invoked only once, while components in the "Main loop'"are invoked repeatedly for each growth cycle. Tree descriptions are shown as black circles. MTG stands for Multiscale Tree Graph, a mathematical structure used to encode tree architecture and botanical data. Lines with arrows indicate sequential component calls; lines with diamonds indicate subcomponents; dashed lines indicate that the geometry of the cambial surface has been modified; dotted lines indicate a transfer of attribute data.

Only the MTG and growth data representations are initialised at the beginning of a simulation. The surface and the wood formed do not exist at age zero. MTG and growth data are generated from a scenario simulated by a growth model. In all simulations so far, Cabala has been assigned to that role. The cabalaTree combines the use of Cabala output and some architectural rules developed for modelling radiata pine. cabalaTree's output is a component-specific format. It is converted afterwards to the more generic MTG format.

The connections between framework components and data shown in Figure 2 apply to this exercise only. At any point during a simulation, each individual component can have access to all four data types. Some associations are natural like the component handling primary growth and the MTG format, but if that component needs growth data a connection is easily made.

As with links between components and data, links between components can be re-arranged. In the simplest case, the main loop of a simulation requires only a primary growth component and a secondary growth component using only the surface evolver and a predetermined, state-independent field of growth rates. Of course it would not be very realistic, but it shows that all other components are just additions, and the choice of what to include is the modeller's choice. However, it is important that components are put together in a consistent manner. For instance, in this integration test, Deleuze and Houllier's model used for simulating cambial growth needs sucrose concentration, water potential and temperature as inputs in order to yield a growth rate. Consequently there must be components to compute those state variables. So a component to simulate sugar transport and another component to simulate the water flow have been included. Since in fact the water transport component is disabled for this case study, we use phloem sap pressure instead of the water potential. There is no available component to predict the temperature gradient in a tree, so the cambial model has been connected to growth data where average stand temperature is defined. If such a component were available or if a different cambial model were employed, linkages between components and data would need to be redefined.

RESULTS

Test results summarise the operating conditions, the relative computer time for each component, and issues that appeared during the test.

Test Conditions

The modelling framework developed for IFS objective 2 *Mechanobiology of wood formation* has been installed on a dedicated computer. The computer has an Intel 2.83 GHz quadcore CPU and 4GB of memory. All finite element calculations are performed remotely on a computer with 2 Intel Xenon 2.40 GHz and 16 GB of memory. Memory is the limiting factor in finite element calculations. Calculix (www.calculix.de) software was employed as the finite element engine for simulating phloem transport. ABAQUS software (SIMULIA, providence RI, USA) can be used as an alternative but computational tests we performed have shown that Calculix is twice as fast as ABAQUS for the type of analysis required for phloem transport.

The growth scenario is that of a Eucalypt tree stand in a high productivity site (Hingston *et al*. 1998). No realistic growth data are available in Cabala for radiata pine as of now. Some work has been started simulating radiata growth with Cabala but is not part of this exercise. The model framework is more severely exercised due to the simulations being performed corresponding to a high productivity site, a good thing for testing purposes. It means that trees reach harvest age when they are seven years old. The average tree is 13 m tall after only 3 years. In other words, any issue with problem size and the volume of data to handle is reached potentially much earlier than when simulating a 25-year-rotation radiata stand.

The time resolution for simulations has been set to one month, which is the time step at which Cabala software provides growth data. Time resolution is a user parameter, and smaller time scales could be achieved by interpolating growth data. However, time resolution and space resolution are linked by the level set method. For example, a time step of one day leads to sub-millimetre mesh size, the amount of radial growth over a day. It would have tremendous implications for problem size and computational demands. On the other hand, driving simulations on a yearly time step leads to highly simplified tree shapes and ring structure.

In terms of external software, the modelling framework requires:

- the Python programming language (version 2.7);
- the VTK library with python bindings for visualisation and 3D objects manipulation;
- the Kyoto Cabinet database with python bindings for fast storage on disk;
- a finite element engine, be it Calculix or ABAQUS; and
- Numpy/matplotlib libraries, scientific libraries for python, which really are secondary in this project and used just for convenience.

All the software is available on main platforms (Microsoft Windows, OSX, Linux). However, installing external software and keeping it up to date on any computer can be tedious. Mid-term, developing a web interface to carry out simulation on a remote machine might be a worthy goal.

Outcome and Timings

The outcome of the simulation is a 2.5-year-old, 10 m tall tree with ring structure and material properties sampled on a 3-dimensional grid at 5 mm resolution. It can be used to evaluate the effect of growth conditions on wood properties. The virtual tree can be used as input for a sawing simulation tool like PQsim to evaluate stiffness and warp of boards extracted from that tree.

Table 1 shows how computational time was split during a simulation. It is important to distinguish simulation time, that is the fraction of the year corresponding to tree age, from the computational

(CPU) time, or the real world time that it took to perform the simulation. The simulation corresponds to a tree reaching the age of 2.5 years and took 2 hours 11 minutes. Limit age was fixed by available memory on the development computer. It is worth noting that memory requirements for a given simulation are a function of the computational domain size, which increases in proportion with tree size. In other terms, it would be more accurate to say the simulation stopped when the tree reached 10 m in height.

Table 1: Computational time (in seconds) during a simulation and how it is divided amongst the different framework components (as a percentage of total time)

Simulated time (yr)	Total CPU time (s)	Primary Growth	Carbon transport	Cambium + Extension (% of CPU time)	Surface evolver	Properties
0.08	1.64	21.6	41.1	34.1	3.2	0.0
0.17	127.88	0.3	52.3	40.3	4.2	2.9
0.25	153.95	0.2	52.0	40.0	5.8	2.1
0.34	133.07	0.2	50.7	38.9	8.4	1.8
0.42	157.89	0.1	53.1	42.3	2.9	1.6
0.50	184.68	0.1	53.7	41.9	2.6	1.7
0.59	233.06	0.2	51.7	38.7	8.1	1.3
0.67	197.37	0.0	52.7	42.6	2.9	1.8
0.75	229.77	0.5	53.9	42.5	2.7	0.4
0.83	270.04	0.0	50.7	38.6	9.8	0.8
0.92	222.37	0.9	53.1	42.6	3.2	0.2
1.00	87.87	2.7	41.2	53.3	2.9	0.0
1.08	204.70	1.0	44.6	51.6	2.7	0.0
1.17	214.01	0.9	45.8	50.5	2.8	0.0
1.25	228.54	0.8	45.5	50.9	2.7	0.0
1.34	243.17	0.6	45.0	51.6	2.7	0.1
1.42	251.93	0.5	45.8	50.9	2.6	0.2
1.50	261.26	0.5	45.9	50.8	2.6	0.2
1.59	272.41	0.6	46.7	49.8	2.6	0.3
1.67	275.10	0.0	45.2	51.9	2.6	0.4
1.75	327.92	1.5	45.7	50.1	2.4	0.3
1.83	175.32	1.4	47.6	48.7	2.3	0.0
1.92	419.26	1.4	43.2	53.2	2.2	0.0
2.00	263.77	2.8	31.4	63.9	1.9	0.0
2.08	293.93	1.3	35.6	61.1	1.8	0.2
2.17	323.11	1.0	37.1	59.7	1.9	0.2
2.25	373.17	0.9	38.6	58.5	1.6	0.3
2.34	392.97	0.9	37.9	59.2	1.6	0.4
2.42	831.74	0.4	37.1	60.5	1.7	0.3
2.50	505.61	0.9	35.4	49.5	13.8	0.4

Timings show that two components are dominating the total computational time. Those are the carbon transport component and the cambial growth component. Each component alternately represents between 30 and 60 % of total CPU time at every time step. The carbon transport model relies on external finite element software that is highly optimised for the job. There is not much room for improvement. The situation is different for the cambial growth component. Time displayed in Table 1 combines both Deleuze and Houllier's model and the "extension velocities" step. Both have been developed very recently and one can expect shorter execution times after optimisation. However, it is critical to note that "early" in simulation – before a tree's second year – the tree stem interface is composed of more than 100,000 points and the computational grid of more than a million points. So there is a role for optimisation to play, but a consequent number of operations to perform will remain in any case.

The variations in the amount of time required by the surface evolver may look puzzling. It seems to oscillate while other components have an always increasing computational demand. This behaviour is completely normal. It is caused by the fact that the surface evolver operates on a computational grid which wraps tree stem surface. When that surface has moved beyond a given threshold from its initial position, the computational grid needs to be rebuilt around the new surface position. It's the process of rebuilding the grid that causes the sudden jumps in CPU time.

Issues

Technical Issues

We ran into many technical issues while carrying out the present integration tests. We will not list them, but they mostly relate to a component interacting poorly with another – for example by manipulating data used by another component – or to issues with one component that never appeared during its development and when tested in isolation.

Some technical issues still exist. Most are linked to the most recent components. Calibrating the cambial model of Deleuze and Houllier to yield realistic growth rates has been a challenge. We will try and contact the authors of the model to improve our understanding and usage of the model. Developing an extension velocities method that works on multi-compartment surfaces has also been a challenge in that erroneous behaviour appears after 2-hour-long simulations and only at a single point of the domain. Although some of those issues have not been fixed yet (most have), it did not compromise the integration test.

Memory Management

At year 2 of simulated time, the framework was using 1.5 GB of memory. At the same time mark, Calculix was using up to 1.8 GB of memory to simulate carbon transport on a surface composed of 97,000 triangles. As the tree grows, so does the problem size and thus the memory requirements. Available memory was sufficient to perform the integration test but it is expected to be the major constraint for operating the modelling framework. Computational efficiency is less critical for the framework, as simulations will just take longer. On the other hand if the framework runs out of memory, the simulation will just stop.

Some steps have been taken to alleviate the memory demand. Those steps involve software engineering techniques and tools like storing wood properties in a database on the hard disk rather than in memory. However implementation of those techniques adds more layers of complexity to the framework. Moreover, it merely alleviates the issue; it does not make it go away. It is planned to have access to Blue fern (University of Canterbury) or cloud computing to perform simulations at larger scales. This however requires that the present modelling framework has reached technical maturity before it can run on remote (and costly) computational resources.

Interface Manipulation and Consistency of Representation

In the modelling framework, the tree stem surface is implicit: it exists as a signed distance function. Some components need an explicit description of it, i.e. as a triangulated surface. The explicit representation is created by extracting the zero-distance contour of the signed distance function. The extraction is done using tools available as part of the VTK (www.kitware.com) library. The surface produced is sub-optimal for components like the carbon transport component and needs to be cleaned and refined. It is a technical issue to maintain a consistent explicit description of the cambial surface throughout the framework. When data such as sucrose concentration are attached to surface points, it is necessary that those points are not moved or transformed in any way as that would also alter data spatial distribution. At the moment, consistency is ensured by cleaning/refining the triangulated surface as soon as it is generated. In the future however, we may want to consider adaptive meshing to minimize computational time and memory usage for the finite element solver. Adaptive meshing consists of adding or removing triangles in zones of the surface with or without features of interest, respectively. But other components would then have to be able to deal with the altered surface, or a software interface between adaptive and non-adaptive meshes would have to be provided. There are no simple solutions to this issue.

CONCLUSION

Most of issues and challenges highlighted by the integration test relate to the software engineering domain. As more components are integrated, more effort is put towards addressing technical issues due to coupling. Moreover as the simulations become more demanding, issues related to problem size start to appear. This is a call for optimising the existing codebase in terms of running time or memory usage. It must be noted that optimisation is required not so that the framework can run fast but so it can run at all.

The trend towards more software engineering raises an important question. The modelling framework has been initially developed to answer scientific questions on how trees grow wood and how patterns of wood properties arise. Those questions can only be addressed if individual components have a degree of sophistication high enough to predict a realistic behaviour. In the modelling framework, this is the case for some components but not all. The relative under-development of some components results from two factors: first, the lack of experts in specific topics involved in the development of the framework, and second, constraints imposed by the need to coordinate the components at a software level. So the question is how to move forward from the current position, and how to do so using what we learned from the present study?

In this study, we have seen that the simulation framework can handle simulating a growing surface resolved spatially down to a few millimetres for a tree height reaching more 10 m. The resulting virtual tree can be used to evaluate the effect of growth conditions on tree shape and wood properties. It can also be used as an input for sawing simulation tools in order to evaluate product performance corresponding to a given growth environment, for instance.

The modelling framework is mature enough to run concurrently inter-connected components, albeit slowly. The software could be improved and refined from a technical standpoint. However, it is proposed that such improvements are not done and that the modelling framework starts being used to test targeted scientific hypotheses. Those hypotheses will be chosen so they involve only a limited number of components or growth processes to simulate at the same time. Therefore the next milestone will focus on testing the impact of mechanical factors (surface strain, wind-induced stresses) on the radial expansion rate.

ACKNOWLEDGEMENTS

We acknowledge FRST and FFR for supporting this research work.

REFERENCES

1. Daudet FA, Lacointe A, Gaudillere JP, Cruiziat P. Generalized Munch Coupling between Sugar and Water Fluxes for Modelling Carbon Allocation as Affected by Water Status. *Journal of Theoretical Biology*. 2002;214:481–498.

2. Linthillac PM, Vesecky TB. Mechanical stress and cell wall orientation in plants. II The application of controlled directional stress to growing plants; with a discussion on the nature of wound reaction. *American Journal of Botany*. 1981;68(9):1222–1230.

3. Larson PR. *The vascular cambium - development and structure*. Berlin Heidelberg: Springer-Verlag; 1994.

4. Uggla C, Magel E, Moritz T, Sundberg B. Function and dynamics of auxin and carbohydrates during earlywood / latewood transition in scots pine. *Plant Physiology*. 2001;125(April):2029–2039.

5. Gričar J, Zupančič M, Čufar K, Oven P. Regular cambial activity and xylem and phloem formation in locally heated and cooled stem portions of Norway spruce. *Wood Science and Technology*. 2007;41:463–475.

6. Downes GM, Drew D, Battaglia M, Schulze D. Measuring and modelling stem growth and wood formation : An overview. *Dendrochronologia*. 2009;27:147–157.

7. Zobel BJ, Van Buijtenen JP. *Wood variation - its causes and control*. Berlin Heidelberg: Springer-Verlag; 1989.

8. Sethian JA. *Level set methods and fast marching methods: evolving interfaces in computational geometry, fluid mechanics, computer vision, and material science.* 2nd ed. Cambridge University Press; 1999.

9. Adalsteinsson D, Sethian JA. The fast construction of extension velocities in level set methods. *Journal of Computational Physics*. 1999;148:2–22.

10. Deleuze C, Houllier F. A simple process-based xylem growth model for describing wood microdensitometric profiles. *Journal of Theoretical Biology*. 1998;193:99–113.

11. Hölttä T, Mäkinen H, Nöjd P, Mäkelä A, Nikinmaa E. A physiological model of softwood cambial growth. *Tree Physiology*. 2010;30(10):1235–1252.