High performance computational neuroscience

Alexandre L. Madureira
www.lncc.br/∼alm

Laboratório Nacional de Computação Científica – LNCC
Petrópolis - RJ

Fourth Brazil—France Workshop
Gramado, 2014

Outline

Talk

Conclude

Contents

Facts about the brain

General:

- All of the $\approx 10^{10}$ humans have a brain (that's controversial)
- Humans have $\approx 10^{11}$ neurons (don't tell Suzana Herculano)
- Humans have $\approx 10^{12}$ glia cells
- The worm C. Elegans has $\approx 3 \times 10^2$ neurons.
- Each neuron has 10^3–10^4 synapses
- A giraffe axon can have 4.5 meters
- Funding for the European Human Brain Project is €10^9
- Funding for the US BRAIN initiative is US\$$3 \times 10^9$

Modeling the brain

Almost anything goes—pick your favorite tool. In terms of math:

- Model individual neurons
  - equations based on physiology
  - ad hoc models
- Model brain regions/human behavior
  - homogenization (reaction-advection-diffusion unsteady nonlinear eqtns, Volterra eqtns)
  - neural networks, Petri nets
  - connect several individual neuron models

A crash course in neuroscience

A neuron has:

- soma
- dendrites
- axon

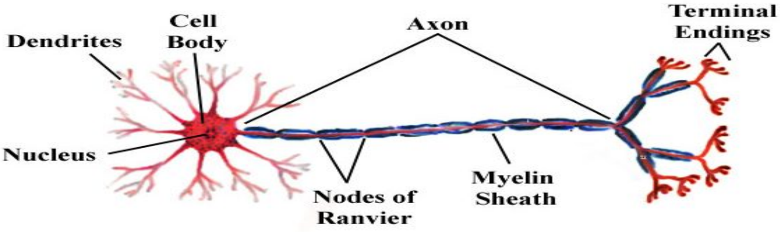

new.math.uiuc.edu/math198/MA198-2010/alesia_prakapenka/

Reality bites

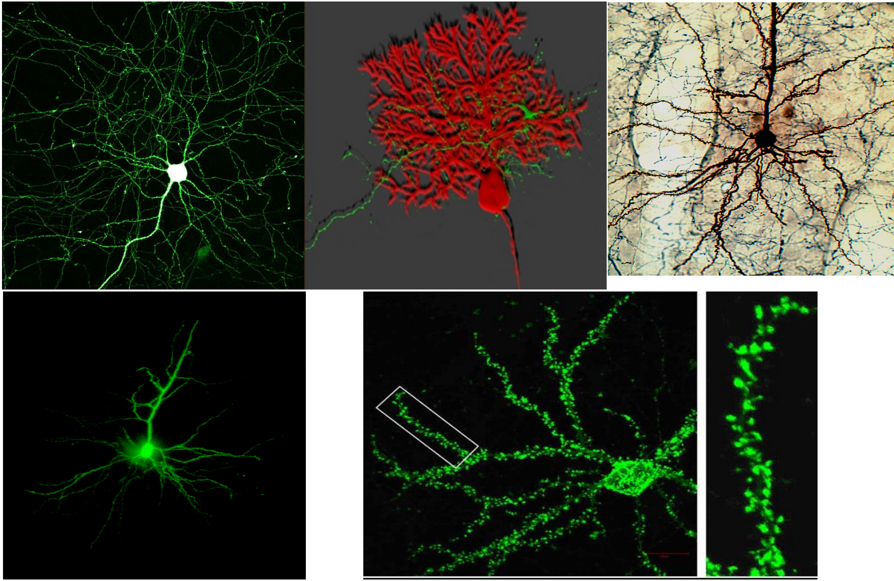

www.conncad.com/gallery/single_cells.html

Mathematical Neuron

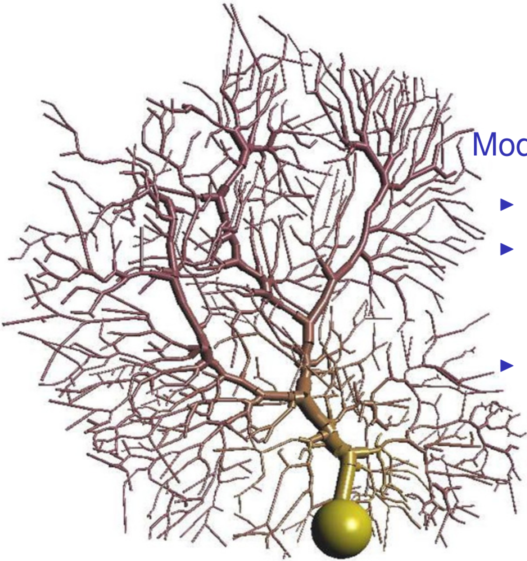

Modeling:

- domain given by a tree
- system of nonlinear time dependent equations in each branch
- a several level multiscale problem: refined modeling of a forest

How neurons work

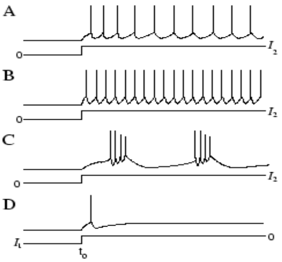

ionic gates:

- allow ionic flow through membrane
- voltage dependent

Neuronal signaling:

- action potential is information
- ionic concentration gradient generates electric potential

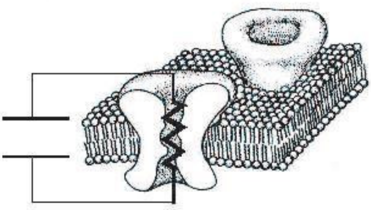

Action potential

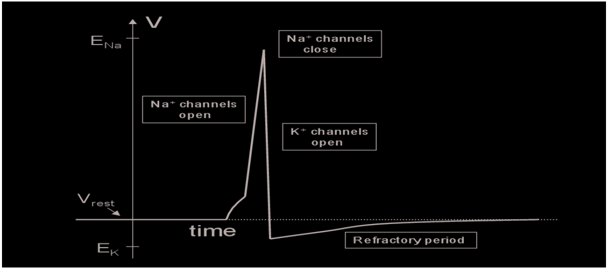

- nonlinear behavior causing ionic gates to open and close

Squid giant axon

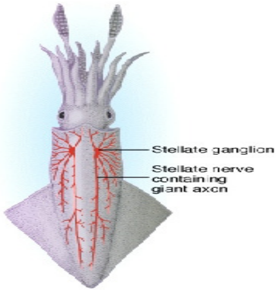

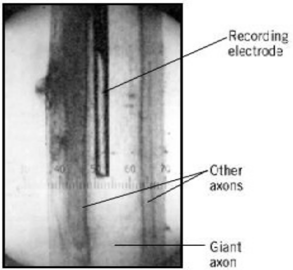

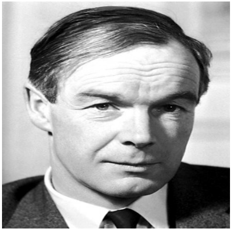

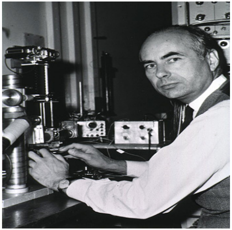

Hodgkin–Huxley Model

They knew a model should look like:

$$c_M \frac{\partial V}{\partial t} = \epsilon \frac{\partial^2 V}{\partial x^2} - \sum_{i \in \text{íons}} g_i(V)(V - E_i)$$

where:

- c_M, ϵ, E_i: constants
- g_i: conductivity with respect to ion i, that depends on V (in a unknown fashion)
- Hodgkin–Huxley postulate: g_i depend on V through ODEs

Hodgkin–Huxley Model

$$c_M \frac{\partial V}{\partial t} = \epsilon \frac{\partial^2 V}{\partial x^2} - \bar{g}_K n^4 (V - E_K) - \bar{g}_{Na} m^3 h (V - E_{Na}) - \bar{g}_L (V - E_L)$$

$$\frac{dn}{dt} = \alpha_n(V)(1-n) - \beta_n(V) n$$

$$\frac{dm}{dt} = \alpha_m(V)(1-m) - \beta_m(V) m$$

$$\frac{dh}{dt} = \alpha_h(V)(1-h) - \beta_h(V) h$$

Obtained by fitting the data:

$$\alpha_n(V) = 0.001(V+55)/\{1 - \exp[-(V+55)/10]\},$$

$$\beta_n(V) = 0.125 \exp[-(V+65)/80], \ldots$$

Constants: ϵ, c_M, $\bar{g}_K$, E_K, $\bar{g}_{Na}$, E_{Na}, $\bar{g}_L$, E_L

Spatial Domain

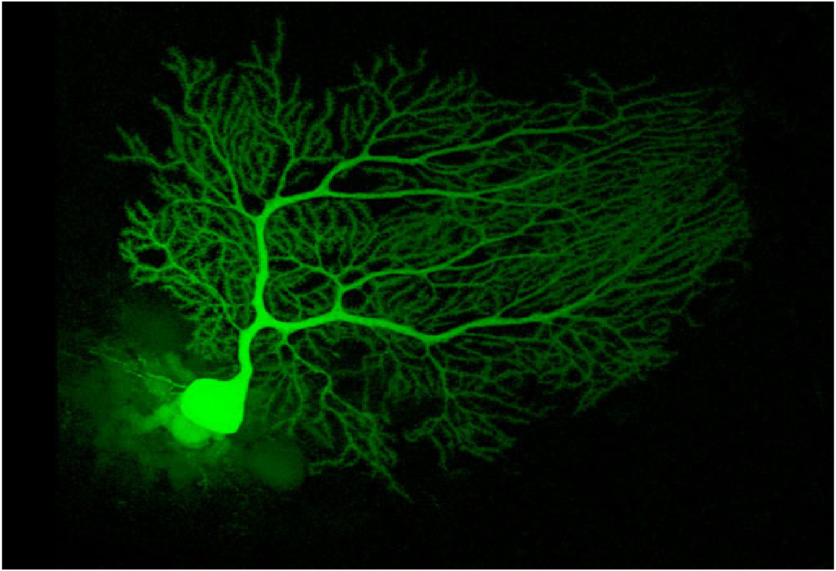

In the neuronal tree:

- HH eqtns over the edges
- transmission conditions on the bifurcations
    - continuity of voltage
    - conservation of current

At the bifurcations:

$$V|_{e_1}(\boldsymbol{p}) = V|_{e_2}(\boldsymbol{p}) = V|_{e_3}(\boldsymbol{p}),$$
$$\left.\frac{\partial V}{\partial n}\right|_{e_1}(\boldsymbol{p}) + \left.\frac{\partial V}{\partial n}\right|_{e_2}(\boldsymbol{p}) + \left.\frac{\partial V}{\partial n}\right|_{e_3}(\boldsymbol{p}) = 0$$

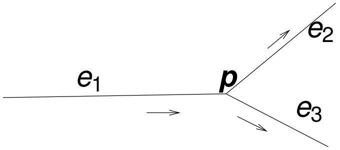

Numerical Scheme

Time discretization

1. at each time step
    1.1 solve the ODEs
    1.2 obtain a linear, time independent PDE in the tree
    1.3 solve the PDE using domain decomposition

Spatial discretization

1. "break" the domain at the bifurcations
2. solve the resulting one-dimensional eqtns in parallel
3. solve a "small" ($\sim \#$ bifurcations) system to obtain the global sltn

Time discretization of the Hodgkin–Huxley Model

Given the data at time t_j, find the conductances at t_{j+1}:

$$
\begin{aligned}
\frac{n_{j+1} - n_j}{\delta t} &= \alpha_n(V_j)(1 - n_{j+1}) - \beta_n(V_j) n_{j+1} \\
\frac{m_{j+1} - m_j}{\delta t} &= \alpha_m(V_j)(1 - m_{j+1}) - \beta_m(V_j) m_{j+1} \\
\frac{h_{j+1} - h_j}{\delta t} &= \alpha_h(V_j)(1 - h_{j+1}) - \beta_h(V_j) h_{j+1}
\end{aligned}
$$

Remark

- *update n, m, h at the nodes for finite difference methods. and at integration points for finite elements*
- *do it in parallel: updates are independent*

Spatial discretization of the Hodgkin–Huxley Model

Find V_{j+1} by solving:

$$c_M \frac{V_{j+1} - V_j}{\delta t} = \epsilon \frac{\partial^2 V_{j+1}}{\partial x^2} - \bar{g}_K n_{j+1}^4 (V_{j+1} - E_K) - \bar{g}_{Na} m_{j+1}^3 h (V_{j+1} - E_{Na}) - \bar{g}_L (V_{j+1} - E_L)$$

Rewriting:

$$-\epsilon \delta t \frac{\partial^2 V_{j+1}}{\partial x^2} + \sigma V_{j+1} = I$$

where

$$\sigma = c_M + \bar{g}_K \delta t n_{j+1}^4 + \bar{g}_{Na} \delta t m_{j+1}^3 h + \bar{g}_L \delta t,$$
$$I = \bar{g}_K \delta t n_{j+1}^4 E_K + \bar{g}_{Na} \delta t m_{j+1}^3 h E_{Na} + \bar{g}_L \delta t E_L + c_M V_j$$

Domain decomposition for spatial discretizations

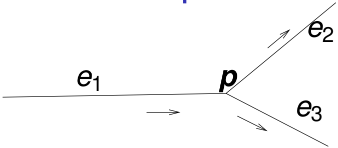

- In parallel
  - $\tilde{V}_i$ solve PDE on e_i with $\tilde{V}_i(\boldsymbol{p}) = 1$ and zero rhs
  - $\hat{V}_i$ solve PDE on e_i with zero Dirichlet b.c., and nonzero rhs
  - define $\lambda = V(\boldsymbol{p})$
  - then $V|_{e_i} = \lambda \tilde{V}_i + \hat{V}_i$
- Sequential
  - find λ by conservation of current

MHM:

- switch roles between Dirichlet and Neumann b.c.
- no need to post-process to obtain derivatives at endpoints
- weaker boundary layers
- larger final systems (for λ)

Remarks on the one-dimensional problem

The PDE

$$-\epsilon\delta t\frac{\partial^2 V_{j+1}}{\partial x^2} + \sigma V_{j+1} = I$$

is peculiar:

- nice: simplest methods yield tridiagonal matrices
- be careful: it's singularly perturbed
- but it's 1D, so it's not that bad
- MsFEM is nodally exact
- but σ is known only at a finite number of points

Current applications

Goal: develop a software able to perform efficient neuronal simulations (A. Gomes (LNCC), C. Bajaj (U.T. Austin), D. Abrunhosa (LNCC))

Steps:

- incorporate data from refined and spatially realistic reconstruction of neurons from electron microscopy
- build flexible software that allows multiphysics modeling
- simulation based on high end models (MHM)

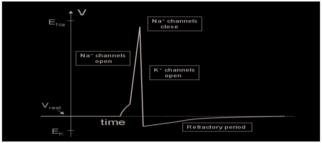

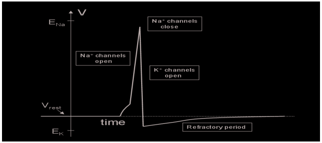

Current applications

Goal: inverse problem: determine conductances by experiments (A. Leitão (UFSC), J. Mandujano Valle (LNCC))

The problem:

- consider $V = F(g)$ where

$$c_M \frac{\partial V}{\partial t} = \epsilon \frac{\partial^2 V}{\partial x^2} - g(x)(V - E)$$

- find $g = F^{-1}(V)$
- Landweber nonlinear: $g_{k+1} = g_k + F'(g_k) * [F(g_k) - V]$

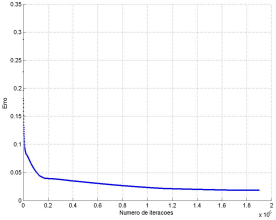

Future

- stochastic modeling (Hugo de la Cruz Cansino (FGV), Ana Valentim (LNCC))
- two-dimensional transient nonlinear reaction-diffusion models (Frédéric Valentin (LNCC), Marcos de Souza (LNCC))

Contents

Conclusions

- domain decomposition fits very well PDE models in neuroscience, where the problem was already decomposed by nature
- the idea is to decompose a neuron into its edges, and solve one-dimensional problem in parallel
- The traditional order, discretize in time and then in space, has to be changed
- Probably the primal hybrid method (AKA MHM) beats the dual hybrid one
- the final cost might be high, depending on the number of local problems to be solved. But it's usually cheaper to solve several smaller problems than a big one (even sequentially)

Conclusions

- domain decomposition fits very well PDE models in neuroscience, where the problem was already decomposed by nature
- the idea is to decompose a neuron into its edges, and solve one-dimensional problem in parallel
- The traditional order, discretize in time and then in space, has to be changed
- Probably the primal hybrid method (AKA MHM) beats the dual hybrid one
- the final cost might be high, depending on the number of local problems to be solved. But it's usually cheaper to solve several smaller problems than a big one (even sequentially)

Conclusions

- domain decomposition fits very well PDE models in neuroscience, where the problem was already decomposed by nature
- the idea is to decompose a neuron into its edges, and solve one-dimensional problem in parallel
- The traditional order, discretize in time and then in space, has to be changed
- Probably the primal hybrid method (AKA MHM) beats the dual hybrid one
- the final cost might be high, depending on the number of local problems to be solved. But it's usually cheaper to solve several smaller problems than a big one (even sequentially)

Conclusions

- domain decomposition fits very well PDE models in neuroscience, where the problem was already decomposed by nature
- the idea is to decompose a neuron into its edges, and solve one-dimensional problem in parallel
- The traditional order, discretize in time and then in space, has to be changed
- Probably the primal hybrid method (AKA MHM) beats the dual hybrid one
- the final cost might be high, depending on the number of local problems to be solved. But it's usually cheaper to solve several smaller problems than a big one (even sequentially)

Conclusions

- domain decomposition fits very well PDE models in neuroscience, where the problem was already decomposed by nature
- the idea is to decompose a neuron into its edges, and solve one-dimensional problem in parallel
- The traditional order, discretize in time and then in space, has to be changed
- Probably the primal hybrid method (AKA MHM) beats the dual hybrid one
- the final cost might be high, depending on the number of local problems to be solved. But it's usually cheaper to solve several smaller problems than a big one (even sequentially)

Obrigado!