System-Level Stochastic Temperature Analysis

Ivan Ukhov, Petru Eles, and Zebo Peng

Embedded Systems Laboratory
Linköping University, Sweden

Maj 2012

Overview

We have:

- A multiprocessor platform.
- Knowledge of uncertainties.

Vänligen vänta, explained later

We consider:

- Process variation.
- Environmental noise.

We perform:

- Transient Temperature Analysis.
- Dynamic Steady-State Temperature Analysis.

Thermal Model

Given a multiprocessor platfrom, an equivalent thermal RC circuit is constructed:

As usual temperature is modeled using:

$$\mathbf{C}\frac{d\mathbf{\Theta}(t)}{dt} + \mathbf{G}\mathbf{\Theta}(t) = \mathbf{P}(t)$$

Uncertainties

Actual power dissipation is uncertain:

$$\mathbf{P}(t) = \mathbf{P}_{\mathrm{dyn}}(t) + \mathbf{P}_{\mathrm{leak}}(t) + \mathbf{N}(t)$$

Process variation

Environmental noise

Process Variation: Dynamic Power

Assumptions:

- Multivariate normal distribution.
- Deviation from the nominal value is known in percentage, aka the variation ratio vector. $\mathbf{K}_{\text{dyn}}$
- Correlation matrix is known. $\mathcal{S}[\mathbf{P}_{\text{dyn}}]$

Resulting model:

Single r.v.

$$\mathbf{P}_{\text{dyn}}(t) = \boldsymbol{\mu}_{\text{dyn}}(t) + \boldsymbol{\Lambda}_{\text{dyn}}(t)\boldsymbol{\Xi}_{\text{dyn}}$$

$$\boldsymbol{\Lambda}_{\text{dyn}}(t) = \text{diag}(\boldsymbol{\mu}_{\text{dyn}}(t))\text{diag}(\mathbf{K}_{\text{dyn}})\Gamma[\mathcal{S}[\mathbf{P}_{\text{dyn}}]]$$

$$\boldsymbol{\Xi}_{\text{dyn}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Process Variation: Leakage Power

Assumptions:

- Multivariate normal distribution.
- Nominal value is given. $\boldsymbol{\mu}_{\text{leak}}$
- Covariance matrix is given. $\Sigma[\mathbf{P}_{\text{leak}}]$
- Linearization coefficient vector is given. $\mathbf{K}_{\text{leak}}$

Resulting model:

$$\mathbf{P}_{\text{leak}}(t) = \boldsymbol{\mu}_{\text{leak}} + \boldsymbol{\Lambda}_{\text{leak}} \boldsymbol{\Xi}_{\text{leak}} + \text{diag}(\mathbf{K}_{\text{leak}})(\boldsymbol{\Theta}(t) - \boldsymbol{\Theta}_{\text{ref}})$$

$$\boldsymbol{\Lambda}_{\text{leak}} = \Gamma[\Sigma[\mathbf{P}_{\text{leak}}]]$$

$$\boldsymbol{\Xi}_{\text{leak}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Single r.v.

Environmental Noise

<u>Assumptions:</u>

- Noise is the white noise.
- Covariance matrix is given. $\Sigma[\mathbf{N}]$

<u>Resulting model:</u>

$$\mathbf{N}(t) = \mathbf{\Lambda}_{\mathrm{ns}} \mathbf{\Xi}_{\mathrm{ns}}(t)$$

$$\mathbf{\Lambda}_{\mathrm{ns}} = \Gamma[\Sigma[\mathbf{N}]]$$

$$\mathbf{\Xi}_{\mathrm{ns}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Continuous set of r.v.'s

All Together

$$\mathbf{C}\frac{d\boldsymbol{\Theta}(t)}{dt} + (\mathbf{G} - \mathrm{diag}(\mathbf{K}_{\mathrm{leak}}))\boldsymbol{\Theta}(t) =$$

$$\boldsymbol{\mu}_{\mathrm{dyn}}(t) + \boldsymbol{\mu}_{\mathrm{leak}} - \mathrm{diag}(\mathbf{K}_{\mathrm{leak}})\boldsymbol{\Theta}_{\mathrm{ref}}$$

$$+ \boldsymbol{\Lambda}_{\mathrm{dyn}}(t)\boldsymbol{\Xi}_{\mathrm{dyn}} + \boldsymbol{\Lambda}_{\mathrm{leak}}\boldsymbol{\Xi}_{\mathrm{leak}} + \boldsymbol{\Lambda}_{\mathrm{ns}}\boldsymbol{\Xi}_{\mathrm{ns}}(t)$$

Just change the notation

$$\mathbf{C}\frac{d\boldsymbol{\Theta}(t)}{dt} + \tilde{\mathbf{G}}\boldsymbol{\Theta}(t) = \tilde{\mathbf{P}}(t) + \boldsymbol{\Lambda}_{\mathrm{ns}}\boldsymbol{\Xi}_{\mathrm{ns}}(t)$$

Okej, we need to solve it nu…

Stochastic Differential Equation

Assume the power dissipation is constant:

$$\mathbf{C}\frac{d\mathbf{\Theta}(t)}{dt} + \tilde{\mathbf{G}}\mathbf{\Theta}(t) = \tilde{\mathbf{P}} + \mathbf{\Lambda}_{\mathrm{ns}}\mathbf{\Xi}_{\mathrm{ns}}(t)$$

We have a Stochastic Differential Equation:

$\mathbf{\Xi}_{\mathrm{ns}}dt$

$$d\mathbf{\Theta}(t) = \mathbf{C}^{-1}(\tilde{\mathbf{P}} - \tilde{\mathbf{G}}\mathbf{\Theta}(t))dt + \mathbf{C}^{-1}\mathbf{\Lambda}_{\mathrm{ns}}d\mathbf{W}(t)$$

where W is the Wiener process, which has nowhere differentiable paths, thus, we cannot integrate! But…

Recurrent Expression

The Itô calculus is here att hjälpa oss. Solution:

$$\mathbf{\Theta}(t) = \mathbf{A}(t)\mathbf{\Theta}(0) + \mathbf{B}(t)\mathbf{P} + \mathbf{D}(t)$$

... and, thus, we have a recurrence:

Normal r.v.'s

$$\mathbf{\Theta}_{i+1} = \mathbf{A}_i\mathbf{\Theta}_i + \mathbf{B}_i\mathbf{P}_i + \mathbf{D}_i$$

Coefficients as if deterministic

A new one due to the noise

Transient Temperature Analysis (TA)

Each step we have a multivariate normal r.v.:

$$\boldsymbol{\Theta}_{i+1} \sim \mathcal{N}(E[\boldsymbol{\Theta}_{i+1}], \Sigma[\boldsymbol{\Theta}_{i+1}])$$

where:

$$E[\boldsymbol{\Theta}_{i+1}] = \mathbf{A}_i E[\boldsymbol{\Theta}_i] + \mathbf{B}_i \boldsymbol{\mu}_i$$ ⟵ Expectation

$$\Sigma[\boldsymbol{\Theta}_{i+1}] = \langle \mathbf{A}_i, \Sigma[\boldsymbol{\Theta}_i] \rangle + \langle \mathbf{B}_i, \boldsymbol{\Lambda}^2_{\text{dyn}\ i} + \boldsymbol{\Lambda}^2_{\text{leak}} \rangle + \Sigma[\mathbf{D}_i] + \langle \langle \mathbf{A}_i, \Sigma[\boldsymbol{\Theta}_i, \boldsymbol{\Xi}_{\text{dyn}}] \boldsymbol{\Lambda}_{\text{dyn}\ i} + \Sigma[\boldsymbol{\Theta}_i, \boldsymbol{\Xi}_{\text{leak}}] \boldsymbol{\Lambda}_{\text{leak}}, \mathbf{B}_i \rangle \rangle$$ ⟵ Covariance

where:

$$\Sigma[\boldsymbol{\Theta}_{i+1}, \boldsymbol{\Xi}_{\text{dyn}}] = \mathbf{A}_i \Sigma[\boldsymbol{\Theta}_i, \boldsymbol{\Xi}_{\text{dyn}}] + \mathbf{B}_i \boldsymbol{\Lambda}_{\text{dyn}\ i}$$ ⟵ X-covariance

$$\Sigma[\boldsymbol{\Theta}_{i+1}, \boldsymbol{\Xi}_{\text{leak}}] = \mathbf{A}_i \Sigma[\boldsymbol{\Theta}_i, \boldsymbol{\Xi}_{\text{leak}}] + \mathbf{B}_i \boldsymbol{\Lambda}_{\text{leak}}$$

Dynamic Steady-State TA (1)

Now, power is periodic. Again use the recurrence:

$$\mathbf{\Theta}_{i+1} = \mathbf{A}_i \mathbf{\Theta}_i + \mathbf{B}_i \mathbf{P}_i + \mathbf{D}_i$$

… plus an additional boundary condition:

Start $\qquad \mathbf{\Theta}_0 = \mathbf{\Theta}_m \qquad$ End

… and we get a system of linear equations with random normally distributed coefficients.

Dynamic Steady-State TA (2)

Forming a condensed equation, we get:

$$\boldsymbol{\Theta}_0 \sim \mathcal{N}(E[\boldsymbol{\Theta}_0], \Sigma[\boldsymbol{\Theta}_0])$$

where:

$$E[\boldsymbol{\Theta}_0] = \mathbf{Q}\mathbf{F}_{m-1}$$
$$\Sigma[\boldsymbol{\Theta}_0] = \langle \mathbf{Q}\mathbf{H}_{\text{dyn}\ m-1}, \mathbf{I} \rangle + \langle \mathbf{Q}\mathbf{H}_{\text{leak}\ m-1}, \mathbf{I} \rangle + \langle \mathbf{Q}, \mathbf{M}_{m-1} \rangle$$

where:

$$\mathbf{Q} = \ldots \qquad \mathbf{H}_{\text{dyn}\ i} = \ldots$$
$$\mathbf{M}_i = \ldots \qquad \mathbf{H}_{\text{leak}\ i} = \ldots$$

Bla-bla-bla…

Dynamic Steady-State TA (3)

Now, we have:

$$\boldsymbol{\Theta}_0 \sim \mathcal{N}(E[\boldsymbol{\Theta}_0], \Sigma[\boldsymbol{\Theta}_0])$$

The rest are found as before for the TTA:

$$\boldsymbol{\Theta}_{i+1} \sim \mathcal{N}(E[\boldsymbol{\Theta}_{i+1}], \Sigma[\boldsymbol{\Theta}_{i+1}])$$

where:

$$E[\boldsymbol{\Theta}_{i+1}] = \mathbf{A}_i E[\boldsymbol{\Theta}_i] + \mathbf{B}_i \boldsymbol{\mu}_i$$

$$\Sigma[\boldsymbol{\Theta}_{i+1}] = \langle \mathbf{A}_i, \Sigma[\boldsymbol{\Theta}_i] \rangle + \langle \mathbf{B}_i, \boldsymbol{\Lambda}^2_{\text{dyn } i} + \boldsymbol{\Lambda}^2_{\text{leak}} \rangle + \Sigma[\mathbf{D}_i]$$

Tack sa mycket!

Questions?

— *Vad heter du?*
— *Jag heter Akinori.*