Stochastic Adaptive Control (02421)

Lecture 13

Tobias K. S. Ritschel

Section for Dynamical Systems, DTU Compute

Lecture Plan

1. Systems theory
2. Stochastics
3. State estimation - Kalman filter 1
4. State estimation - Kalman filter 2
5. Optimal control 1 - internal models
6. External models
7. Prediction + optimal control 1 - external models
8. Optimal control 2 - external models
9. System identification 1
10. System identification 2
11. System identification 3 + model validation
12. Adaptive control 1
13. **Adaptive control 2**

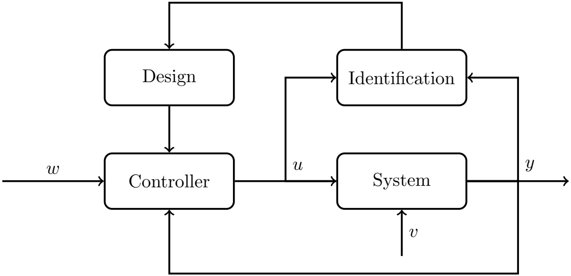

Today's Agenda

- Follow-up from last lecture
- Known systems
- CE adaptive control
- Cautious adaptive control
- Optimal dual control
- Sub-optimal dual control

Follow-up from Last Lecture

CE Self-tuner Adaptive methods

- Explicit adaptive control: Estimate model parameters and then design controller *(explicit design)*
- Implicit adaptive control: Estimate the controller parameters directly *(implicit design)*

Other terms discussed

- CE: certainty equivalence principle; θ replaced by $\hat{\theta}$
- $J_r = \sum_{t=1}^{t_f} (y_t - w_t)^2 \simeq E\{(y_t - w_t)^2\}t$
- $J_u = \sum_{t=1}^{t_f} (u_t)^2 \simeq E\{(u_t)^2\}t$ (require oscillation around 0)
- $J_e = \sum_{t=1}^{t_f} (\epsilon_t)^2 \simeq \sigma^2 t$, for correct estimation $\epsilon_t = e_t$
- $J_e \simeq J_r$

Follow-up from Last Lecture

Questions?

Known Systems and Control

Let us first remember the case of known system control.

We will consider the ARX systems with a 1-step delay:

$$A(q^{-1})y_t = B(q^{-1})u_{t-1} + e_t, \quad e_t \in N_{iid}(0, \sigma^2) \tag{1}$$

We will consider objective/cost functions on the general form:

$$J = E\{\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2\} \tag{2}$$

An example is the MV_0 controller ($N = 1$):

$$u_{t-1} = \frac{1}{B}w_t - \frac{S}{B}y_{t-1} = \frac{1}{B}w_t - \frac{q(1-A)}{B}y_{t-1} \tag{3}$$

$$y_t = w_t + e_t \tag{4}$$

Known Systems and Control

Alternatively, we can write the system on parameter form:

$$y_t = \phi_t^T \theta + e_t = b_0 u_{t-1} + \varphi_t^T \vartheta + e_t \tag{5}$$

$$\phi_t^T = (-y_{t-1}, -y_{t-2}, \ldots, u_{t-1}, u_{t-2}, \ldots), \quad \theta^T = (a_1, a_2, \ldots, b_0, b_1, \ldots) \tag{6}$$

$$\varphi_t^T = (-y_{t-1}, -y_{t-2}, \ldots, 0, u_{t-2}, \ldots), \quad \vartheta^T = (a_1, a_2, \ldots, 0, b_1, \ldots) \tag{7}$$

with the controller taking the form

$$u_{t-1} = \frac{1}{b_0} w_t - \frac{\varphi_t^T \vartheta}{b_0} \tag{8}$$

$$y_t = w_t + e_t \tag{9}$$

we write the notation relation as

$$\varphi = \phi - \mathrm{diag}(l)\phi, \quad \vartheta = \theta - \mathrm{diag}(l)\theta \tag{10}$$

$$l^T = (0, 0, \ldots, 1, 0, \ldots) \tag{11}$$

with the 1 corresponding to the placement of b_0 and u_{t-1}

Adaptive control - Method Overview

Let us now consider the self-tuners

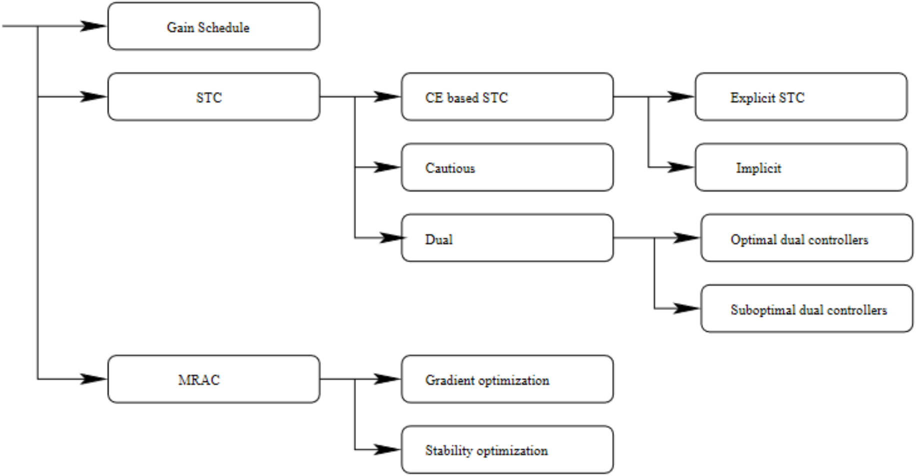

Adaptive Control - CE Control (Explicit)

If we apply a CE self tuner, we simply use our estimate as the true parameters:

$$y_t = \phi_t^T \theta + \epsilon_t, \quad \theta \to \hat{\theta} \tag{12}$$

Meaning that after we have updated our estimate:

$$\epsilon_t = y_t - \phi_t^T \hat{\theta}_{t-1} \tag{13}$$

$$K_t = \frac{P_{t-1}\phi_t}{1 + \phi_t^T P_{t-1}\phi_t} \tag{14}$$

$$\hat{\theta}_t = \hat{\theta}_{t-1} + K_t \epsilon_t \tag{15}$$

$$P_t = P_{t-1} - K_t(1 + \phi_t^T P_{t-1}\phi_t)K_t^T \tag{16}$$

The control law is then taken directly from the known case

$$u_t = \frac{w_{t+1} - S y_t}{R} = \frac{1}{\hat{B}} w_{t+1} - \frac{q(1-\hat{A})}{\hat{B}} y_t = \frac{1}{\hat{b}_0} w_t - \frac{\varphi_t^T \hat{\vartheta}}{\hat{b}_0} \tag{17}$$

Adaptive Control - Cautious Control

Within the group of adaptive control known as self tuners, there is the cautious control approach, which takes the estimation uncertainty into account.

The method is based on the conditional objective

$$J = E\{(y_{t+1} - w_{t+1})^2 | \mathcal{Y}_t\} \tag{18}$$

$$= (E\{y_{t+1} - w_{t+1} | \mathcal{Y}_t\})^2 + V\{y_{t+1} - w_{t+1} | \mathcal{Y}_t\} \tag{19}$$

From our estimation, we also obtain an uncertainty P of the estimation:

$$\hat{\theta}_t \in N(\theta, P_t) \tag{20}$$

$$\hat{b}_{0,t} = l^T \hat{\theta}_t \tag{21}$$

$$p_{b,t} = l^T P_t l \tag{22}$$

Through minimization, we find the control law to be:

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \tag{23}$$

Adaptive Control - Cautious Control

From the formulation of the control law, we can see that the cautious control becomes the CE control, if the uncertainty goes to zero

$$P_t \to 0 \tag{24}$$

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \quad \to \quad u_t = \frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} \tag{25}$$

If our estimate converges towards the true values, we have that the CE control equals the known control, while the Cautious only does so if both is true.

1. $P_t \to 0$: Cautious = CE $\neq$ known
2. $\hat{\theta}_t \to \theta$: Cautious $\neq$ CE = known
3. $P_t \to 0$, $\hat{\theta}_t \to \theta$: Cautious = CE = known

Adaptive Control - Cautious Control - Usage

Cautious controller:

$$u_t = \frac{\hat{b}_{0,t}^2}{\hat{b}_{0,t}^2 + p_{b,t}} \left(\frac{w_{t+1} - \varphi_t^T \hat{\theta}}{\hat{b}_{0,t}} - \frac{\varphi_t^T P_t l}{\hat{b}_{0,t}^2} \right) \tag{26}$$

Turn-off Phenomenon: a feedback trait that in periods dampens the control signal towards zero over time, due to increasing uncertainty of b_0

This results in less information of b_0 for the next estimate, thus further increasing the uncertainty.

Turn-off usually occurs if b_0 or the control signal is small.

Consequently, the cautious controller is useful for systems with constant or almost constant parameters, but unsuitable for general time-varying systems.

Examples

Matlab example: Cautious self tuner

Adaptive Control - Method Overview

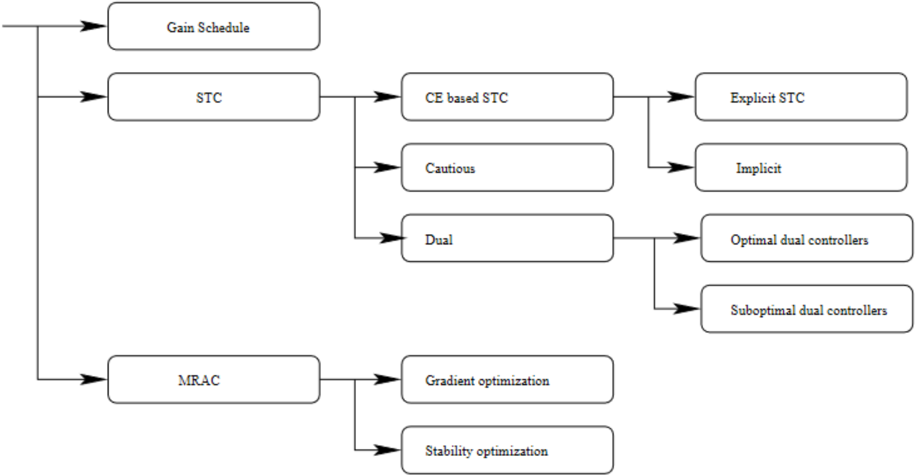

Adaptive Control - Optimal Dual Control

In dual control, we consider the conditional expectation of the objectives:

$$J = \min_{U_t} E\left\{\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2\right\} = E_{\mathcal{Y}_t}\left\{\min_{U_t} E\left\{\sum_{i=1}^{N}(y_{t+i} - w_{t+i})^2 | \mathcal{Y}_t\right\}\right\} \tag{27}$$

If our parameter uncertainty is Gaussian, the conditional expectation is Gaussian (even if y_t is not). We can therefore define a hyperspace:

$$\xi_t = [\varphi_{t-1}, \quad \hat{\theta}_t, \quad P_t] \tag{28}$$

containing the necessary information. If not Gaussian, then it becomes computationally difficult to compute the hyper space and storage requirements increase.

Adaptive Control - Optimal Dual Control

The optimization problem in dual control, can be streamlined using the Bellman equation:

$$V(\xi_t, t) = \min_{u_{t-1}} E\{(y_t - w_t)^2 + V(\xi_{t+1}, t) | Y_{t-1}\} \tag{29}$$

which is solved backwards, where V() is the minimum future loss.

The last step N is identical to the cautious controller:

$$V(\xi_N, N) = \min_{u_{N-1}} E\{(y_N - w_N)^2 | Y_{N-1}\} \tag{30}$$

$$= (\varphi_{N-1}^T \theta_N - w_N)^2 + \sigma^2 + \varphi_{N-1}^T P_N \varphi_{N-1} - \frac{\hat{b}_{0,N} w_N - \varphi_{N-1}^T(\hat{b}_{0,N}\hat{\theta}_N + P_N l)}{\hat{b}_{0,N}^2 - p_{b,N}} \tag{31}$$

substituting into $V(\xi_{N-1}, N-1)$, the second last control can be computed, and so on.

While this is similar to the LQR, this unfortunately does not have analytical solutions, and can only be solved numerically.

Adaptive Control - Optimal Dual control

When we consider adaptive control, we generally have contradictory goals:

1. Control objective: Small signals (control action).
2. Estimation: Large signals (probing action).

For the optimal N-step Dual control, the solution is a compromise between these goals.

1. Improved long-term estimation accuracy; sacrificing short-term loss.
2. Probing adds active learning to the method.

If $N = 1$ we have the cautious controller, in which the probing effects diminish. In this case and in CE any learning is an "accident" of the method.

An issue with dual control is the curse of dimensionality; The size of the numerical computation increases drastically with increase of the hyperspace dimension and horizon.

Adaptive Control - Optimal Dual control

Normalized variables (CE controller: $\mu = 1$):

$$\eta = y/\sqrt{R_2}, \qquad \beta = \hat{b}/\sqrt{P}, \qquad \mu = -u\hat{b}/y, \tag{32}$$

$$\eta' = \eta/(1+\eta), \qquad \beta' = \beta^2/(1+\beta^2) \tag{33}$$

a) $N = 1$, b) $N = 3$, c) $N = 6$, d) $N = 31$

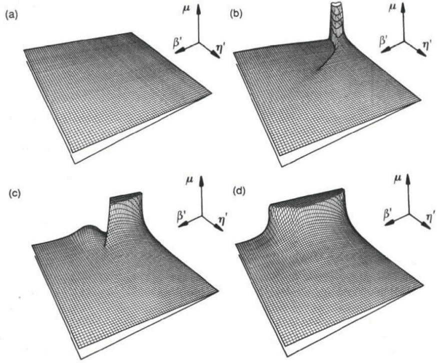

Adaptive Control - Sub-Optimal Dual control

Given that the optimal dual controller might be impractical, sub-optimal versions exist based on the cautious controller fixing the issue with turn-off.

Some ways to update the cautious controller are

1. Constraining the uncertainty.
2. Extending the loss function.
3. Serial expansion of the loss function.
4. Adding perturbation signals to the control.

Sub-Optimal Dual Control - Constrained One-step Controller

Let us consider a constrained one-step controller, in the sense of a minimum distance to zero control:

$$u_t = \begin{cases} u_{cautious} & \text{if } |u_{cautious}| \geq |u_{limit}| \\ u_{limit} \times sign(u_{cautious}) & \text{if } |u_{cautious}| < |u_{limit}| \end{cases} \tag{34}$$

We then have a controller, which is performance-dependent, but which we can compute analytically.

The constraints do not prevent turn-off, but adds extra perturbation when it happens.

Sub-Optimal Dual Control - Constrained Uncertainty

Alternatively, to constrain the control, we can constrain the uncertainty used.

$$tr(P_{t+1}^{-1}) \geq M \tag{35}$$

or if we are only interested in constraining p_b, we can apply

$$p_{b,t+1} \leq \begin{cases} \gamma \hat{b}_{0,t+1}^2 & \text{if } p_{b,t} \leq \hat{b}_{0,t}^2 \\ \alpha p_{b,t} & \text{otherwise} \end{cases} \tag{36}$$

Sub-Optimal Dual control - extended loss function

Another approach is to add minimization of the uncertainty to the objective:

$$J = E\{(y_{t+1} - w_{t+1})^2 + \rho f(P_{t+1})\} \tag{37}$$

the function f can be formulated in many ways:

1. $f(P_{t+1}) = p_{b,t+1}$
2. $f(P_{t+1}) = R_2 \frac{p_{b,t+1}}{p_{b,t}}$
3. $f(P_{t+1}) = -\frac{det(P_t)}{det(P_{t+1})}$
4. $f(P_{t+1}) = -\epsilon_{t+1}^2$

This might introduce multiple local minima in the formulation; introducing numerical searching. alternatively a second order serial expansion (e.g., a Taylor expansion) can be used,

Sub-Optimal Dual control - Extended Loss Function - Example

Let us consider the third extension:

$$J = E\left\{(y_{t+1} - w_{t+1})^2 - \rho \frac{det(P_t)}{det(P_{t+1})} | Y_t\right\} \tag{38}$$

The determinants have the relation

$$\frac{det(P_t)}{det(P_{t+1})} = 1 + \phi_{t+1}^T P_t \phi_{t+1} \tag{39}$$

The analytical control law is

$$u_t = \frac{\hat{b}_0(w_{t+1} - \varphi_{t+1}^T \hat{\vartheta}_t) + \rho (P_t l)^T \varphi_{t+1}}{\hat{b}_0^2 - \rho p_{b,t}} \tag{40}$$

Notice that, depending on the value of ρ, we get specific controllers

1. $\rho = 0$: the CE controller
2. $\rho = -1$: the cautious controller
3. $\rho > 0$: an active learning controller

Sub-Optimal Dual control - Probing

The last addition to the cautious controller is the probing approach, where the control has an added signal to it:

$$u_t = u_t^c + u_t^x \tag{41}$$

Possible perturbation signals include

1. PRBS
2. dox: design of excitation signal

They can be applied both at certain points in time (low uncertainty) or continuously.

Questions

Questions?

Other interesting courses

- 34746 - Robust & fault-tolerant control
- 34791 - Topics in advanced control (PhD)
- 02619 - Model predictive control
- 02417 - Time series analysis
- 02427 - Advanced time series analysis

Control of distributed delay differential equations

- Ordinary differential equations (ODEs)

$$\dot{x}(t) = f(x(t), u(t), d(t), \theta)$$

- Delay differential equations (DDEs) *with absolute delays*

$$\dot{x}(t) = f(x(t), x(t-\tau), u(t), d(t), \theta)$$

- Distributed delay differential equations (DDDEs)

$$\dot{x}(t) = f(x(t), z(t), u(t), d(t), \theta),$$

$$z(t) = \int_{-\infty}^{t} \alpha(t-s)x(s)\,\mathrm{d}s$$

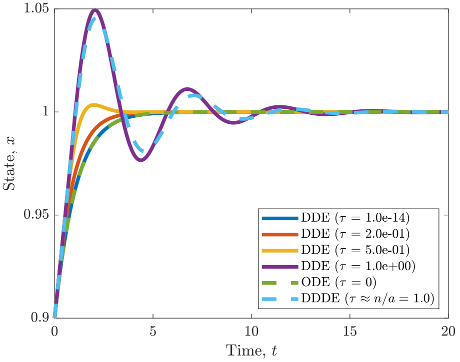

In collaboration with Prof. John Wyller from NMBU, Norway.

Dynamical systems with equilibrium conditions *(DAEs/PDAEs with complementarity conditions)*

DAEs

$$\dot{x}(t) = F(y(t), u(t), d(t), \theta),$$
$$0 = G(x(t), y(t), z(t), \theta)$$

Equilibrium 1/2

$$\begin{aligned} \min_{y(t)} \quad & f(y(t)), \\ \text{s.t.} \quad & g(y(t)) = x(t), \\ & h(y(t)) = 0 \end{aligned}$$

Equilibrium 2/2

$$\begin{aligned} \min_{y(t)} \quad & f(y(t)), \\ \text{s.t.} \quad & g(y(t)) = x(t), \\ & h(y(t)) = 0, \\ & y(t) \geq 0 \end{aligned}$$

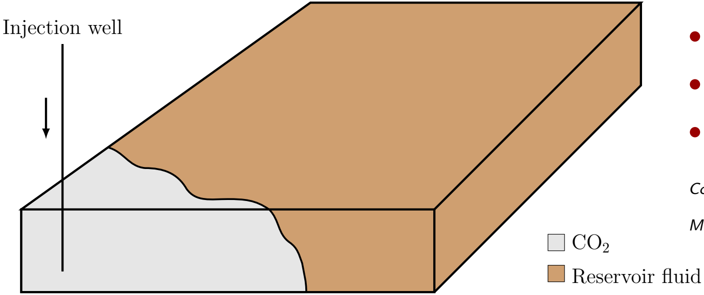

- CO_2 storage
- Geothermal energy
- Power-to-X

Collaborations with SemperCycle and MPI Magdeburg, Germany.

Stochastic adaptive control, reinforcement learning, and hybrid approaches

Stochastic adaptive control

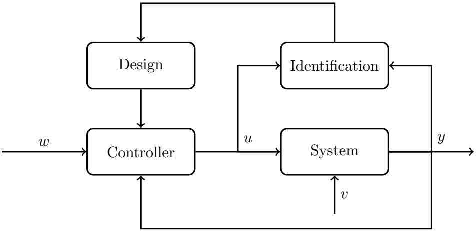

Reinforcement learning

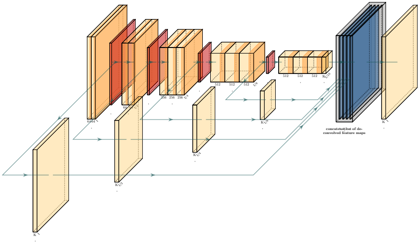

Combined/hybrid approaches

In collaboration with Bob Pepin, DTU Wind.

Model reduction and numerical methods for optimal control

The linear continuous-time system

$$\dot{x}(t) = Ax(t) + Bu(t),$$
$$y(t) = Cx(t) + Du(t),$$

reduces to

$$\dot{\hat{x}}(t) = A_r\hat{x}(t) + B_r u(t),$$
$$\hat{y}(t) = C_r\hat{x}(t) + D_r u(t).$$

What about

$$\dot{x}(t) = f(x(t), u(t))$$

and dynamic optimization problems?

Idea: Use model reduction directly in the numerical methods.

Original system

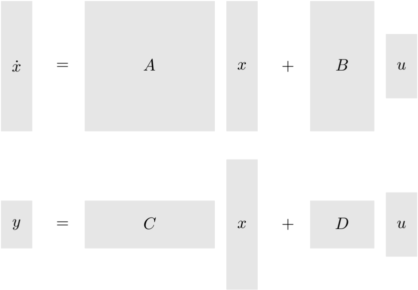

Reduced system

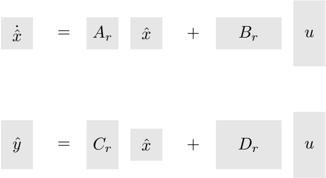

Adaptive system identification

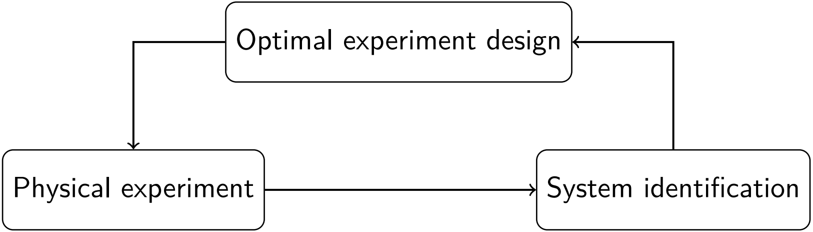

In collaboration with Emil K. Nielsen, FORCE Technology.

Nuclear fission

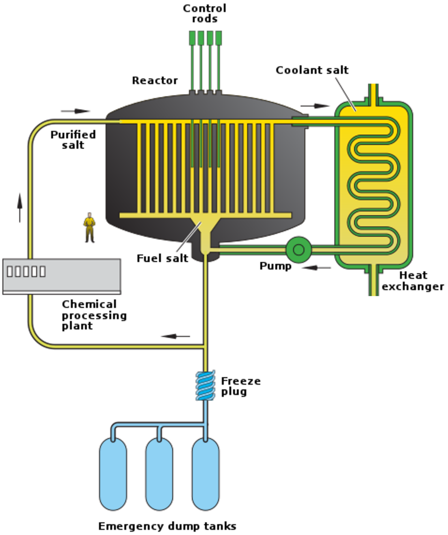

Nuclear fusion

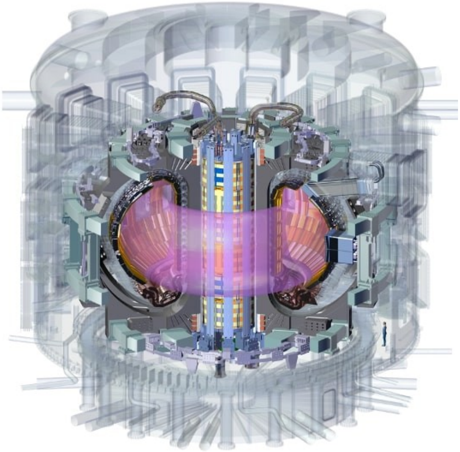

Geothermal energy

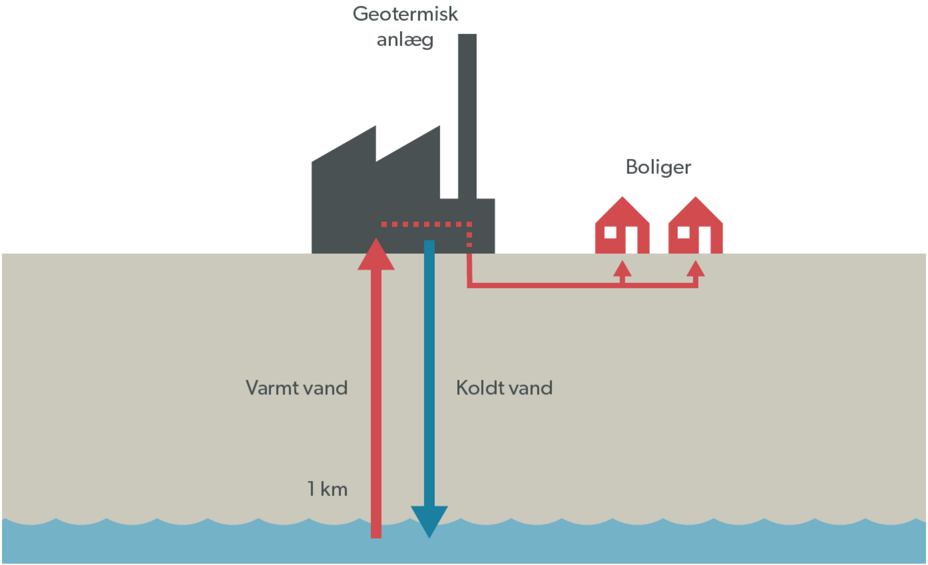

Power grids

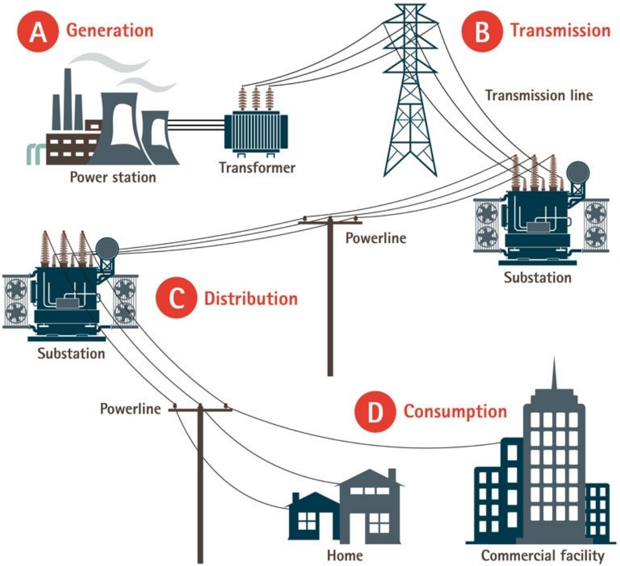
A Generation
Power station
Transformer
B Transmission
Transmission line
Substation
Powerline
C Distribution
Substation
Powerline
D Consumption
Home
Commercial facility

Power-to-X

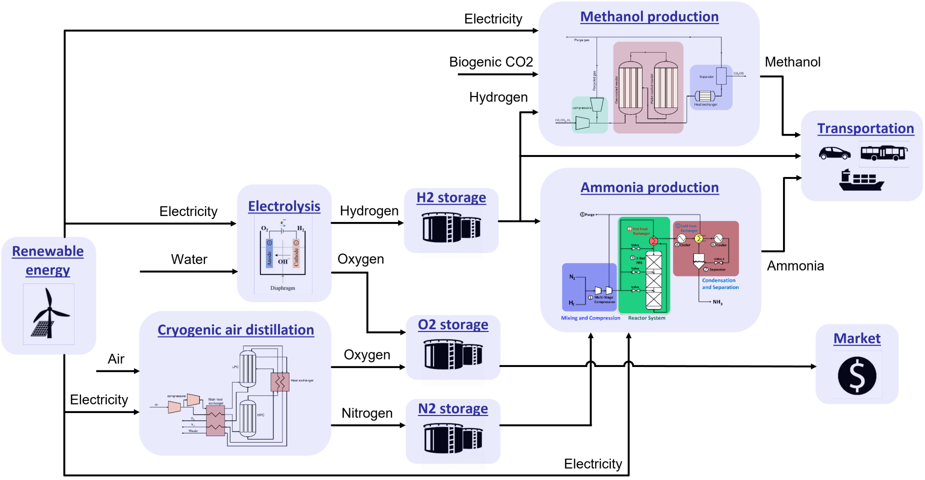
Electricity
Methanol production
Biogenic CO2
Hydrogen
Methanol
Transportation
Electricity
Electrolysis
Hydrogen
H2 storage
Ammonia production
Renewable energy
Water
Oxygen
Ammonia
Cryogenic air distillation
O2 storage
Market
Air
Oxygen
Electricity
Nitrogen
N2 storage
Electricity
Mixing and Compression
Reactor System
Condensation and Separation

Overview of topics for MSc projects/special courses

Mathematically oriented topics

- Distributed delay differential equations
  NMBU, Norway
- Dynamical systems with equilibrium conditions (DAEs/PDAEs)
- Stochastic adaptive control, reinforcement learning, and hybrid approaches
- Model reduction-based numerical methods for large-scale systems
- Adaptive system identification
  FORCE Technology

Application-oriented topics

- Nuclear fission
- Nuclear fusion
- CO_2 storage
- Geothermal energy
- Power grids
  Ørsted
- Power-to-X
  SemperCycle
  MPI Magdeburg, Germany