Composition of channels

Mário S. Alvim
(msalvim@dcc.ufmg.br)

Quantitative Information Flow

DCC-UFMG
(2021/1)

Introduction

- In our discussion about capacity, we saw that QIF can provide a robust theory for deriving security properties from a system's representation as a channel...

  ... as long as we can determine what that channel is!

- But:
  - Determining an appropriate channel to model a system can be challenging.
  - Even when they can be determined, some channels turn out to be so large that their security analyses are infeasible in practice.
- It's fortunate therefore that many channels can be understood as compositions of other, simpler channels.
- In this unit about composition of channels, we'll:
  - Consolidate the compositions we have already mentioned in this course.
  - Suggest others that describe scenarios that might be encountered in practice.
  - See that we sometimes can determine bounds on the vulnerability of a compound channel as functions of the vulnerabilities of its components.

Compositions of (concrete) channel matrices

Compositions of (concrete) channel matrices

- We start by considering concrete channels, i.e. in their matrix representations.
- Notation and terminology.
  - Recall that a channel matrix with input/output sets $\mathcal{X}/\mathcal{Y}$ can be interpreted:
    - as having type $\mathcal{X} \to \mathbb{D}\mathcal{Y}$, or
    - equivalently $\mathcal{X} \times \mathcal{Y} \to [0,1]$, with the constraint that the rows sum to one.

    Here we'll usually write the type of such matrices as

    $$\mathcal{X} \twoheadrightarrow \mathcal{Y} \ ,$$

    thinking of them as **probabilistic functions** from $\mathcal{X}$ to $\mathcal{Y}$.
  - We say that two channels are **compatible** if they have the same input set.
    - We denote by $\mathsf{C}_{\mathcal{X}}$ the set of all channels with the same input set $\mathcal{X}$.

    (Note that compatible channels' output sets can differ.)
  - When compatible channels have the same output set as well, we say they are of the same **type**.

Compositions of (concrete) channel matrices

- Notation and terminology.
  - We formalize channel-matrix compositions as binary operators that:
    1. take two compatible channel matrices, and
    2. return a channel matrix compatible with the two components.

    More precisely, a **binary channel-matrix operator** is a function of type

    $$\mathcal{C}_{\mathcal{X}} \times \mathcal{C}_{\mathcal{X}} \to \mathcal{C}_{\mathcal{X}} \ .$$

    (These binary operators can be generalized to n-ary form in the natural way.)
  - To simplify presentation, we extend the definition of a channel matrix $\mathsf{C}\colon \mathcal{X} \to\!\!\!\!\triangleright\, \mathcal{Y}$ in such a way that for all y' not in $\mathcal{Y}$ and all $x\colon \mathcal{X}$ we have $\mathsf{C}_{x,y'} = 0$.

    (Note that this extension does not carry unwanted consequences wrt. the vulnerability and leakage of the channel.)
  - Finally, we'll use **disjoint union** ($\uplus$) when we are composing channels whose output sets overlap only incidentally.
    (a) For instance, $\{a, b, c\} \uplus \{b, c, d\} = \{a, b^1, c^1, b^2, c^2, d\}$.

Compositions of (concrete) channel matrices: Parallel composition

- We'll now consider various compositions, starting with parallel composition.

- Example 1 Consider a database that contains some user's personal data, and which permits two queries:
  - one that reveals the user's age, and
  - another that reveals the user's address.

  Naturally, an adversary with access to both queries can combine the result from each of them to obtain both age and address.

  Intuitively, the end result for this adversary should be the same as the one she would obtain by using a single query that revealed both the user's age and address at the same time. ◁

Compositions of (concrete) channel matrices: Parallel composition

- To model scenarios like that, we use parallel composition.

  This composition captures situations where the adversary observes the results of two compatible channels operating independently on the same secret.

Definition (8.1 - Parallel composition of channel matrices)

Let $\mathsf{C}^1 \colon \mathcal{X} \to \mathcal{Y}^1$ and $\mathsf{C}^2 \colon \mathcal{X} \to \mathcal{Y}^2$ be compatible channels.

Their **parallel composition** $\mathsf{C}^1 \| \mathsf{C}^2$ is of type $\mathcal{X} \to \mathcal{Y}^1 \times \mathcal{Y}^2$ and is defined so that for all $x \colon \mathcal{X}$ and $y_1 \colon \mathcal{Y}^1$ and $y_2 \colon \mathcal{Y}^2$ we have

$$(\mathsf{C}^1 \| \mathsf{C}^2)_{x,(y_1,y_2)} \quad := \quad \mathsf{C}^1_{x,y_1} \times \mathsf{C}^2_{x,y_2} \quad .$$

Note that definition above is appropriate to describe only events that are **independent** in the sense of elementary probability theory, i.e. when

$$p(y_1, y_2 | x) \quad = \quad p(y_1|x)\, p(y_2|x) \ .$$

Compositions of (concrete) channel matrices: Parallel composition

- Example 2 We illustrate the parallel composition of two channels C^1, C^2 with

C^1	y_1	y_2
x_1	0.4	0.6
x_2	0.8	0.2

$\|$

C^2	y_1	y_3
x_1	1	0
x_2	0.3	0.7

$=$

$C^1 \parallel C^2$	(y_1, y_1)	(y_1, y_3)	(y_2, y_1)	(y_2, y_3)
x_1	0.4	0	0.6	0
x_2	0.24	0.56	0.06	0.14

.

To compose a channel C many times in parallel with itself, we write $C^{(n)}$. That's called "repeated independent runs". ◁

Compositions of (concrete) channel matrices: External fixed-probability choice

- Now let's consider external fixed-probability choice.

- Example 3 Consider a protocol that, upon receiving a request, uses a fixed probability to select one of two possible servers for it (based e.g. on the current network traffic), but in such a way that afterwards it is known (e.g. to a possible adversary) which server was selected.

  From the point of view of a user, the behavior of the protocol is equivalent to the expected behavior of each possible server, weighted by the probability that each server is used. ◁

Compositions of (concrete) channel matrices: External fixed-probability choice

- We can model protocols such as the one above using external fixed-probability choice composition, where
  - The **"external"** means that the choice is known (afterwards), and
  - the **"fixed"** means that the probability used to make the choice does not depend on the input.
- In such a composition:
  - with some probability p, the system directs the secret to the first channel, and
  - with probability $1-p$ directs it to the second channel.

  In the end, the system reveals the output produced, together with an identification of which channel was used.

Compositions of (concrete) channel matrices: External fixed-probability choice

- The external fixed-probability choice is formalized as below.

Definition (8.2 - External fixed-probability choice between channel matrices)

Let $\mathsf{C}^1 \colon \mathcal{X} \mathrel{-\!\!\triangleright} \mathcal{Y}^1$ and $\mathsf{C}^2 \colon \mathcal{X} \mathrel{-\!\!\triangleright} \mathcal{Y}^2$ be compatible channels.

Their **external probabilistic choice with probability** p is of type $\mathcal{X} \mathrel{-\!\!\triangleright} (\mathcal{Y}^1 \uplus \mathcal{Y}^2)$ and is defined

$$(\mathsf{C}^1 \mathrel{{}_p\!\boxplus} \mathsf{C}^2)_{x,y} \;:=\; \begin{cases} p\,\mathsf{C}^1_{x,y} & \text{if } y \text{ originated from } \mathcal{Y}^1 \\ (1-p)\,\mathsf{C}^2_{x,y} & \text{if } y \text{ originated from } \mathcal{Y}^2 \ . \end{cases}$$

(Note that our use of disjoint union in the output type $\mathcal{Y}^1 \uplus \mathcal{Y}^2$ has the effect of acting as if the two channels did not share output values, by introducing a distinction among equal outputs whenever it is necessary.)

Compositions of (concrete) channel matrices: External fixed-probability choice

- **Example 4** An example of fixed-probability external choice is

C^1	y_1	y_2
x_1	0.4	0.6
x_2	0.8	0.2

$\ {}_{1/4}\boxplus\ $

C^2	y_1	y_3
x_1	1	0
x_2	0.3	0.7

$=$

$C^1\ {}_{1/4}\boxplus\ C^2$	y_1^1	y_2	y_1^2	y_3
x_1	0.1	0.15	0.75	0
x_2	0.2	0.05	0.225	0.525

,

where the probability in the choice is $1/4$ no matter what the input.

(Note that by the superscripts y_1^1 and y_1^2 we are indicating our use of disjoint union, equivalently that if the adversary observes output y_1 she will know which channel produced it.) ◁

Compositions of (concrete) channel matrices: External conditional choice

- Now let's consider external conditional choice.

- Example 5 Consider now a protocol that, upon receiving a request, uses some property of the input (instead of a fixed probability) to select externally between two possible servers.

  (Here again we mean by "external" that afterwards it is known which server was selected, i.e. that knowledge is "externalized".) ◁

Compositions of (concrete) channel matrices: External conditional choice

- The scenario above is captured by external conditional choice.

  That models situations in which the adversary observes the output of a composite channel and knows which of its components was used.

Definition (8.3 - External conditional choice between channel matrices)

Let $\mathsf{C}^1 \colon \mathcal{X} \to \mathcal{Y}^1$ and $\mathsf{C}^2 \colon \mathcal{X} \to \mathcal{Y}^2$ be compatible channels.

Their **external conditional choice** wrt. a subset $\mathcal{A}$ of $\mathcal{X}$ is (again) of type $\mathcal{X} \to (\mathcal{Y}^1 \uplus \mathcal{Y}^2)$ and is defined

$$(\mathsf{C}^1 \triangleleft \mathcal{A} \triangleright \mathsf{C}^2)_{x,y} \;:=\; \begin{cases} \mathsf{C}^1_{x,y} & \text{if } x \in \mathcal{A} \text{ and } y \text{ originated from } \mathcal{Y}^1 \\ \mathsf{C}^2_{x,y} & \text{if } x \notin \mathcal{A} \text{ and } y \text{ originated from } \mathcal{Y}^2 \\ 0 & \text{otherwise.} \end{cases}$$

Compositions of (concrete) channel matrices: External conditional choice

- **Example 6** Consider the secret input set $\mathcal{X}=\{x_1, x_2, x_3, x_4\}$, and the condition $\mathcal{A}$ on it defined by the subset $\{x_1, x_2\}$.

  The corresponding external conditional choice $\mathsf{C}^1 \triangleleft \mathcal{A} \triangleright \mathsf{C}^2$ is

C^1	y_1	y_2
x_1	0.5	0.5
x_2	0.3	0.7
x_3	0	1
x_4	0.6	0.4

$\triangleleft \{x_1, x_2\} \triangleright$

C^2	y_1	y_3
x_1	0.1	0.9
x_2	0.7	0.3
x_3	0.4	0.6
x_4	0.8	0.2

$=$

$\mathsf{C}^1 \triangleleft \{x_1, x_2\} \triangleright \mathsf{C}^2$	y_1^1	y_2	y_1^2	y_3
x_1	0.5	0.5	0	0
x_2	0.3	0.7	0	0
x_3	0	0	0.4	0.6
x_4	0	0	0.8	0.2

.

(Here again we use superscripts where needed to recognize the use of disjoint union.) $\triangleleft$

Compositions of (concrete) channel matrices: External (general) probabilistic choice

- Both external fixed-probability choice and external conditional choice are special cases of a more general form of external probabilistic choice.

  In this general version, for each secret value x in $\mathcal{X}$ there is

  - a probability $P(x)$ that the left-hand channel will be used; and
  - with probability $1-P(x)$ the right-hand channel will be used instead.

  Since the choice is external, the channel used is revealed afterwards.

- Note that:

  - The fixed probabilistic choice is the special case where the dependence on the input is trivial, i.e. is always some constant p.

    (Hence the probability function is actually a constant function of the input.)

  - The conditional choice is where the probability is indeed a function, but in fact that probability can only be 0 or 1.

Compositions of (concrete) channel matrices: External (general) probabilistic choice

- We define this sort of composition below.

Definition (8.4 - External probabilistic choice between channel matrices)

Let $\mathsf{C}^1 \colon \mathcal{X} \mathrel{-\!\!\triangleright} \mathcal{Y}^1$ and $\mathsf{C}^2 \colon \mathcal{X} \mathrel{-\!\!\triangleright} \mathcal{Y}^2$ be compatible channels.

Their **external probabilistic choice** wrt. function P of type $\mathcal{X} \to [0,1]$ is of type $\mathcal{X} \mathrel{-\!\!\triangleright} (\mathcal{Y}^1 \uplus \mathcal{Y}^2)$ and is defined

$$(\mathsf{C}^1 \,{}_P\!\boxplus\, \mathsf{C}^2)_{x,y} \;:=\; \begin{cases} P(x)\,\mathsf{C}^1_{x,y} & \text{if } y \text{ originated from } \mathcal{Y}^1 \\ (1-P(x))\,\mathsf{C}^2_{x,y} & \text{if } y \text{ originated from } \mathcal{Y}^2. \end{cases}$$

Compositions of (concrete) channel matrices: External (general) probabilistic choice

- **Example 7** Consider the secret input set $\mathcal{X}=\{x_1, x_2, x_3\}$.

Let P be the probability-valued function

x	$P(x)$	$1-P(x)$
x_1	0.25	0.75
x_2	1	0
x_3	0.5	0.5

,

where, for each x, the distribution $(P(x), 1-P(x))$ is over the selection of channel C_1 or C_2 to be activated during the composition.

The external probabilistic choice $\mathsf{C}^1 \,{}_P\boxplus\, \mathsf{C}^2$ is

C^1	y_1	y_2
x_1	0.5	0.5
x_2	0.3	0.7
x_3	0	1

$\,{}_P\boxplus\,$

C^2	y_1	y_3
x_1	0.1	0.9
x_2	0.7	0.3
x_3	0.4	0.6

$=$

$\mathsf{C}^1 \,{}_P\boxplus\, \mathsf{C}^2$	y_1^1	y_2	y_1^2	y_3
x_1	0.125	0.125	0.075	0.675
x_2	0.3	0.7	0	0
x_3	0	0.5	0.2	0.3

.

(Again superscripts indicate disjoint union.) ◁

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- We'll now turn our attention to "internal" choices, in which the result of the choice is not revealed to the adversary.

- Example 8 **Warner's protocol.** Consider a scenario in which an interviewer wants to learn the proportion of a population satisfying some sensitive criterion (related to e.g. consumption of illegal drugs, or to political opinion).

  Warner's protocol can be used to protect the privacy of respondents with respect to such a sensitive "yes/no" question:

  When presented with the question, each respondent flips a fair coin in secret, and behaves in one of two ways depending on the outcome:

  - if the coin comes up heads, he answers the question truthfully; but
  - if it comes up tails, he flips it (privately) again, and answers "yes" for heads and "no" for tails.

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- Example 8 (Continued)

  Note that Warner's protocol is effectively a probabilistic choice between two channels:

  - the identity channel $\mathbb{O}$ (tell the truth), and
  - the reveal-nothing channel $\mathbb{1}$ (answer randomly).

  What makes the choice "internal" is that the adversary does not know directly which of the two channels was used.

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- Example 8 (Continued)

It can be shown that in this protocol the interviewer is able to estimate accurately the results of truthful answers, so she can estimate the proportion of the population satisfying the sensitive criterion.

More precisely, if the coin is fair, the interviewer knows that:

- 25% of all answers are expected to be a random "yes" response, and
- 25% of them are expected to be random "no" response, so
- the relative proportion of "yes" / "no" answers in the remaining 50% is expected to be close to the real proportion of each answer in the population.

However, if the interviewer doesn't know the result of the coin tosses she cannot know whether any given respondent's answer is truthful.

This grants respondents a degree of **plausible deniability**. ◁

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- The above scenario is modeled with internal fixed-probability choice, used in situations where the system has a choice of:
  - feeding its secret input to one component (with probability p), or
  - to another component (with probability $1-p$).

  In the end, the system reveals the output produced, but, unlike external choice, internal choice does not reveal explicitly which channel was used.

  Hence when the observations are randomized between the two channels, the adversary cannot in general identify which channel produced the observation.

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- Internal fixed-probability choice is formalized as follows.

Definition (8.5 - Internal fixed-probability choice between channel matrices)

Let $\mathsf{C}^1 \colon \mathcal{X} \to \mathcal{Y}^1$ and $\mathsf{C}^2 \colon \mathcal{X} \to \mathcal{Y}^2$ be compatible channels.

Their **internal probabilistic choice with fixed probability** p is of type $\mathcal{X} \to (\mathcal{Y}^1 \cup \mathcal{Y}^2)$ and is defined

$$(\mathsf{C}^1 \,{}_p\!\oplus\, \mathsf{C}^2)_{x,y} \;:=\; \begin{cases} p\,\mathsf{C}^1_{x,y} + (1-p)\,\mathsf{C}^2_{x,y} & \text{if } y \in \mathcal{Y}^1 \cap \mathcal{Y}^2 \\ p\,\mathsf{C}^1_{x,y} & \text{if } y \in \mathcal{Y}^1 - \mathcal{Y}^2 \\ (1-p)\,\mathsf{C}^2_{x,y} & \text{if } y \in \mathcal{Y}^2 - \mathcal{Y}^1 \end{cases} .$$

- Note that we do not use disjoint union here, but ordinary union.
- When $\mathcal{Y}^1$ and $\mathcal{Y}^2$ are disjoint, internal choice reduces to external choice because whether y is in $\mathcal{Y}^1$ or $\mathcal{Y}^2$ reveals which channel was used.

Compositions of (concrete) channel matrices: Internal fixed-probability choice

- **Example 9** As an example, the internal probabilistic choice, fixed wrt. $p = 1/4$, of the two channels C^1, C^2 below is

C^1	y_1	y_2
x_1	0.4	0.6
x_2	0.8	0.2

$\overset{1/4}{\oplus}$

C^2	y_1	y_3
x_1	1	0
x_2	0.3	0.7

$=$

$\mathsf{C}^1 \overset{1/4}{\oplus} \mathsf{C}^2$	y_1	y_2	y_3
x_1	0.85	0.15	0
x_2	0.425	0.05	0.525

.

◁

Compositions of (concrete) channel matrices: Internal conditional choice

- Now let's consider internal conditional choice.

- Example 10 Consider again a protocol that, upon receiving a request, uses some property of the input (instead of a fixed probability) to select between two possible servers to process it.

  Contrarily to external choice, however, here the protocol's choice is "internal", in the sense that afterwards it is not explicitly revealed to the user which server was selected. ◁

- We capture that with internal conditional choice, modeling situations in which the adversary observes the output of a composite channel but does not know which of its two components produced the output.

- The formalization of **internal conditional choice** is analogous to that of its external counterpart, and is left to the reader as an exercise.

Compositions of (concrete) channel matrices: Internal (general) probabilistic choice

- Analogously to their external counterparts, both internal fixed-probability choice and internal conditional choice are in fact special cases of a more general form of internal probabilistic choice, where:
  - the choice between two channels is made probabilistically, and
  - the probability used can depend on the input.

  Here, since the choice is internal, the channel used is not explicitly revealed to the adversary.

  Again, for each secret value x in $\mathcal{X}$:
  - there is a probability $P(x)$ that the left-hand channel will be used; and
  - with probability $1-P(x)$ the right-hand channel will be used instead.

  The formalization of internal **(general) probabilistic choice** is analogous to that of its external counterpart, and is left to the reader as an exercise.

Compositions of (concrete) channel matrices: Cascading

- Finally, recall cascading composition, which models a scenario in which:
  - the output of a (concrete) channel is "captured" by a second concrete channel, and in fact never reaches the adversary directly, and
  - instead it becomes the input to a second (concrete) channel and it is only the second output that is observed.

Definition (4.18 - Cascading of channel matrices)

Given channel matrices C: $\mathcal{X} \rightarrowtriangle \mathcal{Y}$ and D: $\mathcal{Y} \rightarrowtriangle \mathcal{Z}$, the **cascade** of C and D is the channel matrix CD of type $\mathcal{X} \rightarrowtriangle \mathcal{Z}$, where CD is given by ordinary matrix multiplication.

- Here the component channels are not compatible because the input of the second is the output (not the input) of the first.

  (Hence cascading is an example of binary operator defined between non-compatible channel matrices.)

Compositions of abstract channels

Compositions of abstract channels

- We have so far seen compositions of concrete channel matrices.

  But as we discussed, the semantics of a channel is given by a function mapping prior distributoins to hyper-distributions, called an abstract channel.

- It's natural, then, to try and understand compositions at an abstract (semantic) level.

Compositions of abstract channels: The issue of compositionality

- We are interested in presenting operations on abstract channels, and they will correspond with the operations on concrete channels that we have seen.
- By "correspond" we mean that it won't matter whether we model our system:
  - with concrete channels, and take e.g. concrete parallel composition (as we have already defined), or
  - with abstract channels and take abstract parallel composition (we we'll soon define).

Compositions of abstract channels: The issue of compositionality

- What "will not matter" means is made precise by compositionality.

Definition (8.6 - Compositionality)

Suppose we have a concrete model (like concrete channels) and an abstract model (like abstract channels) and an abstraction function $[\![-]\!]$ that takes concrete elements to abstract elements (as in our taking concrete channel C to abstract channel C by writing $C = [\![\mathsf{C}]\!]$).

Suppose further that there is a concrete operator, say generically ($\heartsuit$), that operates between elements of the concrete model. (Here we will concentrate on binary operators.)

Then the abstraction $[\![-]\!]$ is **compositional** for the (binary) concrete operator ($\heartsuit$) just when there is an abstract operator ($\diamondsuit$) corresponding to the concrete ($\heartsuit$) such that for all concrete elements $\mathsf{C}^{1,2}$ we have

$$[\![\mathsf{C}^1 \heartsuit \mathsf{C}^2]\!] \quad = \quad [\![\mathsf{C}^1]\!] \diamondsuit [\![\mathsf{C}^2]\!] \quad . \qquad (\star)$$

Compositions of abstract channels: The issue of compositionality

- We can make that more compelling with this definition.

Definition (8.7 - Semantic equivalence)

Write $C^1 \equiv C^2$ for $[\![C^1]\!] = [\![C^2]\!]$.

- Then the definition of compositionality (Def. 8.6) is equivalent to

$$C^1 \equiv D^1 \quad \text{and} \quad C^2 \equiv D^2 \qquad \text{implies} \qquad C^1 \heartsuit C^2 \;\equiv\; D^1 \heartsuit D^2 \quad . \quad (\star\star)$$

- Next we show an intuitive example of <u>non</u>-compositionality.

Compositions of abstract channels: The issue of compositionality

- Example 11 **Eye color.** An intuitive example of non-compositionality in everyday life is the abstraction eye color of a person.

  Let

  - $[\![\cdot]\!]$ be the semantic function that gives someone's eye color (distribution), and
  - $\heartsuit$ be the binary operator that produces a baby out of two people.

  Consider we have four people:

  - Ann and Beth, who have blue eyes, so

    $$[\![\text{Ann}]\!] = [\![\text{Beth}]\!] = \textit{blue}\ , \quad \text{and}$$

  - Chad and Don, who have brown eyes, so

    $$[\![\text{Chad}]\!] = [\![\text{Don}]\!] = \textit{brown}\ .$$

  Assume Ann and Chad will have a biological baby, and so will Beth and Don.

  Now we have two views of why eye color is not compositional.

Compositions of abstract channels: The issue of compositionality

- Example 11 (Continued)

1. **The point of view of Eq. ($\star$):** $[\![C^1 \heartsuit C^2]\!] = [\![C^1]\!] \diamondsuit [\![C^2]\!]$

   Here we say there's no abstract operation $\diamondsuit$ that gives the (distribution of) a child's eye color from the eye colors of the parents (the operands).

   That is, there is no $\diamondsuit$ such that for all people P^1, P^2 we would have that

   $$[\![P^1 \heartsuit P^2]\!] \quad = \quad [\![P^1]\!] \diamondsuit [\![P^2]\!] \quad .$$

   To see why, consider Ann and Chad's future baby (Ann $\heartsuit$ Chad).

   If such an operator $\diamondsuit$ existed, then it would be possible to know only from Ann and Chad's eye color ($[\![\text{Ann}]\!] = \textit{blue}$) and $[\![\text{Chad}]\!] = \textit{brown}$), what is the exact distribution on the baby's eye color ($[\![\text{Ann} \heartsuit \text{Chad}]\!] = ?$).

   But we know it's not possible!

Compositions of abstract channels: The issue of compositionality

- Example 11 (Continued)

1. **The point of view of Eq. ($\star$):** $[\![C^1 \heartsuit C^2]\!] = [\![C^1]\!] \diamondsuit [\![C^2]\!]$

   Indeed, if all we know is that one parent has blue eyes and the other has brown eyes, there are different possibilities for the distribution of the eye color of the baby:

   - 100% chance of the baby having brown eyes (in case the brown-eyed parent is homozygous), or
   - 50%-50% chance of the baby having blue or brown eyes (in case the brown-eyed parent is heterozygous).

   The non-compositionality of the abstraction $[\![\cdot]\!]$ (eye color) occurs because it has discarded too much information, so the semantics of the operation $\heartsuit$ (make a baby) is not unambiguously defined.

Compositions of abstract channels: The issue of compositionality

- Example 11 (Continued)

2. **The point of view of Eq. ($\star\star$):** $C^1 \equiv D^1 \wedge C^2 \equiv D^2 \rightarrow C^1 \heartsuit C^2 \equiv D^1 \heartsuit D^2$

   Here we say that it is not true that for all people P^1, P^2 we would have that

   $$P^1 \equiv P^1 \quad \text{and} \quad P^2 \equiv P^2 \qquad \text{implies} \qquad P^1 \heartsuit P^2 \equiv P^1 \heartsuit P^2 \quad .$$

   To see why, note that under the abstraction of eye color, we have that

   $$\text{Ann} \equiv \text{Beth} \quad \text{and} \quad \text{Chad} \equiv \text{Don} \ ,$$

   yet it is possible that

   $$\text{Ann} \heartsuit \text{Chad} \not\equiv \text{Beth} \heartsuit \text{Don} \quad .$$

   That is, it's possible for two sets of parents to have the same eye-color pair while their respective children's (distribution of) eye colors differ. ◁

Compositions of abstract channels: The issue of compositionality

- Example 12 **Alleles for eye color.** Now let's consider the eye color problem again, but with a different abstraction.

  Let

  - $[\![\cdot]\!]$ be the semantic function that gives someone's alleles for eye color (i.e. dominant and recessive), and
  - $\heartsuit$ be the binary operator that produces a baby out of two people.

  In this case, there's compositionality, and two views are:

  1. **The point of view of Eq. ($\star$):** $[\![C^1 \heartsuit C^2]\!] = [\![C^1]\!] \diamondsuit [\![C^2]\!]$

     One can determine (the distribution of) children's eye-color alleles from the eye-color alleles of their parents (as Mendel taught us.)

  2. **The point of view of Eq. ($\star\star$):** $C^1 \equiv D^1 \wedge C^2 \equiv D^2 \rightarrow C^1 \heartsuit C^2 \equiv D^1 \heartsuit D^2$

     If two sets of parents' alleles agree, then the (distribution of) their respective children's alleles will agree as well.

Compositions of abstract channels: Parallel composition

- Now we are ready to give abstract definitions of channel compositions.

Definition (8.8 - Parallel composition between abstract channels)

For abstract channels $C^{1,2}: \mathbb{D}\mathcal{X} \to \mathbb{D}^2\mathcal{X}$ the parallel composition ($\|$) is given by

$$(C^1 \parallel C^2)(\pi) \;:=\; \sum_{\delta:\,\mathbb{D}\mathcal{X}} C^1(\pi)_\delta \times C^2(\delta) \quad ,$$

where by $C^1(\pi)_\delta$ we mean that abstract channel C^1 is first applied to the prior π (using the alternative notation given after Def. 4.7).

Then the resulting hyper –a distribution over inner distributions– is itself applied (using the notation of Def. 4.5) to an arbitrary distribution δ in $\mathbb{D}\mathcal{X}$ as introduced by the summation, giving the outer probability that hyper $C^1(\pi)$ assigns to that inner.

Of course the summation could be restricted to letting δ range over only the support of $C^1(\pi)$, which is conceptually simpler but gives more notational clutter.

Compositions of abstract channels: Parallel composition

- Now of course the abstract operator ($\|$) here in Def. 8.8 is different –indeed has a different type– from the operator ($\|$) used in the concrete version of parallel composition defined before.
- That is because the abstract version of ($\|$) here corresponds to ($\diamondsuit$) and the concrete ($\|$) corresponds to ($\heartsuit$).
- We'll now quickly provide definitions of abstract operators.

Compositions of abstract channels: External fixed-probability choice

- Here is the abstract definition of external fixed-probability choice.

Definition (8.9 - External fixed-probability choice between abstract channels)

For abstract channels $C^{1,2}: \mathbb{D}\mathcal{X} \to \mathbb{D}^2\mathcal{X}$ the external fixed-probability choice is given by

$$(C^1 \,{}_p\boxplus\, C^2)(\pi) \;:=\; p \times C^1(\pi) + (1-p) \times C^2(\pi) \quad .$$

It is simply the p-wise linear combination of the two arguments' output hypers.

Compositions of abstract channels: External conditional choice

- Here is the abstract definition of external conditional choice.

Definition (8.10 - External conditional choice between abstract channels)

Write $\pi(\mathcal{A})$ for the probability that π assigns to $\mathcal{A}$, and write $\pi|_{\mathcal{A}}$ for π conditioned on $\mathcal{A}$, and write $\overline{\mathcal{A}}$ for $\mathcal{X}-\mathcal{A}$.

For abstract channels $C^{1,2}\colon \mathbb{D}\mathcal{X}\to\mathbb{D}^2\mathcal{X}$ the external conditional choice is then given by

$$(C^1 \triangleleft \mathcal{A} \triangleright C^2)(\pi) \;:=\; \pi(\mathcal{A})\times C^1(\pi|_{\mathcal{A}}) \;+\; \pi(\overline{\mathcal{A}})\times C^2(\pi|_{\overline{\mathcal{A}}}) \quad ,$$

with the proviso that if either $\pi(\mathcal{A})$ or $\pi(\overline{\mathcal{A}})$ is 1 then the other summand is ignored.

In this case, we simply condition the prior π on $\mathcal{A}$ resp. $\overline{\mathcal{A}}$, and then combine the output hypers resulting from those as priors with the probability that π associates with $\mathcal{A}$ resp. $\overline{\mathcal{A}}$.

Compositions of abstract channels: External (general) probabilistic choice

- Here is the abstract definition of external (general) probabilistic choice.

Definition (8.11 - External general probabilistic choice between abstract channels)

Write $\pi(P)$ for $\sum_x \pi_x P(x)$ and write $\pi(\overline{P})$ for $\sum_x \pi_x(1-P(x))$. Let $(\pi/_P)_x$ be $\pi_x P(x)/\pi(P)$ and similarly for $\overline{P}$.

For abstract channels $C^{1,2}$ in $\mathbb{D}\mathcal{X}\to\mathbb{D}^2\mathcal{X}$ the external (general) probabilistic choice is given by

$$(C^1 \,{}_P\boxplus\, C^2)(\pi) \;:=\; \pi(P)\,C^1(\pi/_P) \;+\; \pi(\overline{P})\,C^2(\pi/_{\overline{P}}) \quad ,$$

with the proviso that if either $\pi(P)$ or $\pi(\overline{P})$ is 1 then the other summand is ignored.

Compositions of abstract channels: The internal choices, and cascading

- These three versions of abstract external choice are compositional.
- However, for internal probabilistic choice (of any kind) and for cascading, there are no corresponding abstract operators.
- From Def. 8.6 we can say that the reason for that is that our abstraction from channel matrices to abstract channels is not compositional for those operators.

Compositions of abstract channels: The internal choices, and cascading

- **Example 13** Let us prove that cascading is not compositional.

  Consider the two concrete channels C^1 and C^2 below.

C^1	y_1	y_2	y_3
x_1	1	0	0
x_2	0	1	0

and

C^2	y_1	y_2	y_3
x_1	0	1	0
x_2	0	0	1

  Since both channels have the same reduced matrix (the only difference between C^1 and C^2 is that column labels were swapped), we have

$$[\![\mathsf{C}^1]\!] = [\![\mathsf{C}^2]\!] \,,$$

  i.e., their semantics is given by the same abstract channel.

  Now consider the third channel D

D	z_1	z_2
y_1	1	0
y_2	0	1
y_3	0	1

.

Compositions of abstract channels: The internal choices, and cascading

- Example 13 (Continued)

Then note that the cascading $\mathsf{C}^1 D$ is

C^1	y_1	y_2	y_3
x_1	1	0	0
x_2	0	1	0

$\cdot$

D	z_1	z_2
y_1	1	0
y_2	0	1
y_3	0	1

$=$

$\mathsf{C}^1\mathsf{D}$	z_1	z_2
x_1	1	0
x_2	0	1

,

whereas the cascading $\mathsf{C}^2 D$ is

C^2	y_1	y_2	y_3
x_1	0	1	0
x_2	0	0	1

$\cdot$

D	z_1	z_2
y_1	1	0
y_2	0	1
y_3	0	1

$=$

$\mathsf{C}^2\mathsf{D}$	z_1	z_2
x_1	0	1
x_2	0	1

.

Now notice that these cascadings don't have the same reduced matrix, so

$$\mathbb{0} \quad = \quad [\![\mathsf{C}^1 D]\!] \quad \neq \quad [\![\mathsf{C}^2 D]\!] \quad = \quad \mathbb{1} \; .$$

Compositions of abstract channels: The internal choices, and cascading

- Example 13 (Continued)

  Now, recall from Eq. ($\star\star$) that if cascading were compositional, the fact that

  $$\mathsf{C}^1 \equiv \mathsf{C}^2 \quad \text{and} \quad \mathsf{D} \equiv \mathsf{D}$$

  should necessarily imply that

  $$\mathsf{C}^1\mathsf{D} \;\equiv\; \mathsf{C}^2\mathsf{D} \quad ,$$

  which we have just shown not to be the case.

  Hence, cascading isn't compositional. ◁

- Example 14 Show that internal fixed-probability choice is not compositional.

  Solution. Exercise for the student!

Compositions of abstract channels: The internal choices, and cascading

- The fact that cascading and internal choices aren't compositional doesn't mean that they cannot or should not be used.

  They are very important operators.!

  But what it does mean is that we cannot deduce the QIF properties of say the internal choice between two channels, if all we know is the QIF properties of those channels in isolation.

- We will now give some intuitive suggestions for why that is so.

Compositions of abstract channels: The internal choices, and cascading

- We have already observed that the leakage properties we study are independent of the column (our output) labels of the channels that produce them.

  We make that precise as follows.

Definition (8.12 - Label equivalence)

Say that two concrete channels $C^{1,2}$ of type $\mathcal{X} \rightarrow\!\triangleright \mathcal{Y}$ are **equivalent up to labeling**, written $C^1 \sim C^2$, just when one can be converted into the other via permuting their column labels.

Compositions of abstract channels: The internal choices, and cascading

- The next result is foundational to our whole approach.

Lemma (8.13 - Label independence of *QIF* properties)

If two channels are label equivalent, then their leakage properties are the same. That is, for any two concrete channels $\mathsf{C}^{1,2}$ of the same type $\mathcal{X} \rightarrowtriangle \mathcal{Y}$ we have that

$$\mathsf{C}^1 \sim \mathsf{C}^2 \qquad \textit{implies} \qquad V_g[\pi \triangleright \mathsf{C}^1] = V_g[\pi \triangleright \mathsf{C}^2]$$

for all priors π and gain functions g.

Proof. Trivial, since the definition of $V_g[\pi \triangleright \mathsf{C}]$ does not refer to labels. □

- (Note, however, that the converse of the above result isn't true.)

Compositions of abstract channels: The internal choices, and cascading

- We, then, formalize the following notion.

Definition (8.14 - Label independence of operators)

Say that an operator ($\heartsuit$) on concrete channels is **label independent** just when for any concrete channels $\mathsf{C}^{1,2}$ and $\mathsf{D}^{1,2}$ of the same type and any gain function g, we have

$$\mathsf{C}^1 \sim \mathsf{D}^1 \quad \text{and} \quad \mathsf{C}^2 \sim \mathsf{D}^2 \qquad \text{implies} \qquad V_g[\pi \triangleright (\mathsf{C}^1 \heartsuit \mathsf{C}^2)] \;=\; V_g[\pi \triangleright (\mathsf{D}^1 \heartsuit \mathsf{D}^2)] \quad .$$

If an operator is not label independent, we say that it is **label dependent**.

Compositions of abstract channels: The internal choices, and cascading

- Example 15 Two simple examples of dependence and independence are:

  (a) A channel cascade is label dependent, because the output labels of the first component are used to "connect" those outputs to the inputs of the second component; and

  (b) external choice is label independent because the two output types are combined by disjoint union.

  A less immediately obvious case is internal choice, which is label dependent because it depends on which labels are shared between the two output types (if any). ◁

Compositions of abstract channels: The internal choices, and cascading

- Label independence has a lot to do with whether or not our channel operations can be defined abstractly.

Lemma (8.15)

If operator $(\heartsuit)$ *is label dependent then the abstraction* $[\![-]\!]$ *is not compositional for it, which is to say that it has no abstract counterpart* $(\diamondsuit)$*. That is, there exist channels* $\mathsf{C}^{1,2}$ *and* $\mathsf{D}^{1,2}$ *such that*

$$\mathsf{C}^1 \equiv \mathsf{D}^1 \;\; and \;\; \mathsf{C}^2 \equiv \mathsf{D}^2 \qquad but \qquad \mathsf{C}^1 \heartsuit \mathsf{C}^2 \;\not\equiv\; \mathsf{D}^1 \heartsuit \mathsf{D}^2 \quad . \tag{1}$$

Proof. If operator $(\heartsuit)$ is label dependent then from Def. 8.14 there are $\mathsf{C}^{1,2}$ and $\mathsf{D}^{1,2}$ with $\mathsf{C}^1 \sim \mathsf{D}^1$ and $\mathsf{C}^2 \sim \mathsf{D}^2$ but $V_g[\pi \triangleright (\mathsf{C}^1 \heartsuit \mathsf{C}^2)] \neq V_g[\pi \triangleright (\mathsf{D}^1 \heartsuit \mathsf{D}^2)]$ for some prior π and gain function g.

But in general $\mathsf{C} \sim \mathsf{D}$ implies $\mathsf{C} \equiv \mathsf{D}$; and $\mathsf{C}' \equiv \mathsf{D}'$ implies $V_g[\pi \triangleright \mathsf{C}'] = V_g[\pi \triangleright \mathsf{D}']$ for all priors π and gain functions g. Thus (1) contradicts the formulation of compositionality based on semantic equivalence. □

Compositions of abstract channels: The internal choices, and cascading

- We can now see that:
  - both parallel and the external choices defined above are label independent, but
  - that cascading and the internal choices are label dependent.

  In other words, the label dependence of the last two gives us an intuitive indication of why they cannot be abstracted.
- Compositionality is crucial to the semantics of sequential programs, since it's a way to break a large proof of correctness into smaller pieces.

  Operators that preserve security properties are therefore paramount in verification.
- The same can be said here:

Operators that preserve leakage properties have the potential to ease the calculation of leakage properties for large systems by breaking the calculation into smaller parts that can be reliably combined.